



CorreLog®

Security Correlation Server User Reference Manual

<http://www.correlog.com> <mailto:info@correlog.com>

CorreLog, User Reference Manual

Copyright © 2008 – 2015, CorreLog, Inc. All rights reserved.

No part of this manual shall be reproduced without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibilities for errors or omissions. Nor is any liability assumed for damages resulting from the use of this information contained herein.

Table of Contents

Section 1: Introduction	5
Section 2: CorreLog Installation	13
Section 3: CorreLog Basic Usage	23
Section 4: CorreLog Applications	39
Section 5: CorreLog Correlation	81
Section 6: CorreLog Automation	99
Section 7: CorreLog Ticketing	119
Section 8: CorreLog Reporting	137
Appendix A: List of All Screens	161
Appendix B: Syslog Protocol	179
Appendix C: Dashboard Gadget List	191
Appendix D: Windows Process List	199
Alphabetical Index	205

Section 1: Introduction

This document is a comprehensive manual on the CorreLog Security Correlation Server, which is a compact software system that aggregates, catalogs, and correlates enterprise-wide system messages, with special application in security and log file monitoring. The CorreLog Server collects and correlates Syslog and other types messages. Additionally, the system processes messages from other locations and in different formats, including Microsoft Windows event log data.

The CorreLog software collects message information in real time, and provides a central repository of these messages on a single server. Using this program, administrators and operators can view aggregate message information from many different network devices. The program provides a secure web-based approach to gathering message information, and provides tools to understand that information, with special application in security management.

This manual should be of interest to network managers and administrators, responsible for installing and maintaining CorreLog. This manual will also be useful to any developers who are interested in using this software as a basis for larger enterprise management strategies. Finally, this information will be essential to security officers of your organization.

The CorreLog Correlation Server

CorreLog is a correlation server. The program resides on a platform in your enterprise, and continuously listens for Syslog messages. As these messages are received, they are logged, and cataloged into related groups of messages. The user can search this information, and can take automatic action on correlated messages, such as updating database, sending notification, or running arbitrary user programs and batch files.

Within your enterprise, CorreLog can have various possible roles:

- **Security Monitoring.** The CorreLog system collects messages, and permits a wide assortment of message filtering and overrides. The system then provides tools to correlate these messages into significant event on the network, especially related to system security.
- **Syslog Receiver.** The CorreLog system collects Syslog messages. The program transparently supports a variety of different message types. Many sites may use CorreLog simply as a Syslog receiver, whereas other sites may use the program mainly as a security and correlation tool.
- **Stand Alone Manager.** The CorreLog system can operate as a stand-alone management tool, collecting data and providing notifications to users when significant events occur. In particular, CorreLog is highly useful as a specialized manager for enterprise-wide log information.
- **Front End Other Network Managers.** The CorreLog system operates as a front-end for other network managers, and can feed these managers correlated messages, suitable for use with reporting facilities. For example, CorreLog can provide the security-monitoring portion of a larger management strategy.
- **Data Archiving and Reporting.** The CorreLog system can operate as a data archiving program, storing large quantities of message information in compressed format, for auditing and forensic purposes. The CorreLog system also generates, distributes, and archives reports.
- **Unattended Data Collection.** CorreLog can operate in a complete "unattended" manner, collecting data without any requirement on the part of users. The program incorporates features to trim log files, and incorporates an "Auto-learn" function that will automatically adjust thresholds and parameters based upon the data it collects.

The CorreLog system is entirely web-based. It does not rely on a console for operation. All configuration activities (except for initial installation) are performed by authorized users via the web interface. The CorreLog system is intended to be

a highly secure way of viewing information, and incorporates various encryption mechanisms, as well as permitting secure HTTP access. The program includes role-based permissions, and allows "admin", "user" and "guest" type logins to be created.

Introduction To Syslog Protocol

The main messaging protocol for the CorreLog system is the "Syslog" protocol. This is a simple, highly interoperable, and well-established standard. (If you are unfamiliar with this messaging protocol, see Appendix B of this document.)

The management of Syslog messages is a valuable, but often-overlooked aspect of network and business management. Within your enterprise, your routers, servers, workstations, and business applications are constantly collecting important error and status information. This status information resides in error logs, transaction logs, and event logs on each computer.

If you are operating a network of any modest complexity, you are already heavily vested in the Syslog protocol. This capability is already built into many of your critical systems. And this information is highly pertinent to your business operations. Syslog protocol is one of the oldest management protocols, therefore is sometimes (incorrectly) perceived as being deprecated or obsolete, when it is actually a mainstay of network management that has become one of the most interoperable protocols currently in existence.

Syslog Support On Win32 Platforms

Syslog is extensively supported on Unix, and is also supported by Cisco and other hardware vendors. However, Syslog is not a supported "out-of-box" function of Windows platforms. To permit CorreLog to work with Windows platforms, the standard CorreLog program includes redistributable utilities. These utilities are referred to as the CorreLog Windows Tool Set (CorreLog WTS), which has its own user manual comprising a part of this program's documentation. The CorreLog WTS can be downloaded from a variety of locations, including the "Home" screen of the CorreLog Server.

- **CorreLog Syslog Message Service Program** The CorreLog Server includes the redistributable "CO-sysmsg.exe" program, which can be installed as a service on Windows platforms. This program monitors the Windows Event Logs, and sends Syslog messages of appropriate facilities and severities to the CorreLog Server.
- **CorreLog Log File Monitor.** The "CO-sysmsg.exe" program, described above, in addition to monitoring the Windows event logs, can be configured continuously "tail" streaming log files, looking for match

patterns. This provides a simple way of instrumenting any Windows streaming log file (including Oracle error logs, HTTP transfer logs, and other application software logs) to send Syslog messages.

- **Sendlog Program.** The CorreLog Server includes the redistributable “sendlog.exe” program, which can send arbitrary Syslog messages to the CorreLog. This utility is easily installed and interfaced with Windows application programs, including the Windows Performance Manager.

Each of the above utilities runs as a non-intrusive Windows service on all Microsoft platforms, including Win200X, XP, and Vista platforms. The user can select any or all of the above techniques to quickly add Syslog capability to a Windows system. See the “CorreLog Windows Tool Set User Reference” manual for more information on these utilities.

Introduction To Correlation

The word "correlation" has various meanings and different interpretations. The most basic definition is that correlation is simply a relationship between two or more things. The relationship can be parallel, causal, reciprocal, linear, or nonlinear, and be associated with functions of time or other functions.

CorreLog performs a "semantic" correlation. The output of the CorreLog Server (i.e. the dependent variable) is a specific "meaning" associated with messages that have been received, or a specific action that is executed by the program.

Operationally, the CorreLog Server finds "meaning" in the messages, through use of simple or complex match patterns that divide messages into "Threads". Additionally, CorreLog employs "Triggers" to establish context to messages, and "Alerts" to monitor specific message rates. Once the meaning of a message (or group of messages) has been determined, CorreLog takes specific action such as sending a Syslog message, running a program, or opening a ticket and assigning this ticket to a user or group.

CorreLog Server Features

- **High Speed Message Reception.** The CorreLog Server is suitable to operate as the single Syslog receiver for all devices on the network of large enterprises. CorreLog can process more than 2000 messages per second and can handle burst traffic of more than 10.000 messages. CorreLog physically tracks and catalogs network devices with out a hard upper limit. It can receive messages from virtually unlimited numbers of sources without tracking.

- **Cataloging, Correlation, Reporting Of Information.** The CorreLog Server provides a powerful correlation service. This includes implementation of correlation "Threads", "Alerts", "Triggers", and "Actions", and "Catalogs. These functions allow the user to reduce and analyze real-time data, and take action. The correlation features require minimal configuration, and serve as building blocks for larger correlation strategies.
- **Large Scale Data Aggregation, Archiving, Reporting Ability.** The CorreLog Server is designed to have high data aggregation ability. It can collect in excess of 1 Gigabyte worth of data each day, and save this data for up to 500 days online, and for more than 10 years (5,000 days) in offline, compressed format. The archiving function includes MD5 checksums and security codes on data items, to support detailed forensics. Reports are generated daily, in Microsoft Excel format.
- **Large Scale Data Searching Ability.** One of the most important functions of the CorreLog system is its search engine capability. CorreLog employs a high speed, real time index system. This allows quick searches through massive amounts of data. Users can search a terabyte of data for a particular keyword in less than one second.
- **Automatic Trimming Of Data.** CorreLog is intended to be maintenance free. Background programs will remove extraneous files, drop the latest data from the system, archive data to other disks, and allow CorreLog to operate with no periodic maintenance scheduling required. These maintenance parameters are settable via the Web interface. CorreLog can easily operate in an "unattended" mode.
- **Data Filtering and Reduction.** CorreLog filters input data by device, facility, severity, or message keyword, or any combination of these. The CorreLog server additionally provides effective message de-duplication functions. This provides a large degree of control over what programs and facilities can add messages to the CorreLog Server message log. Filtered data can be retained as part of the CorreLog archive of messages, useful for forensics.
- **Ability To Define New Syslog Facilities.** CorreLog allows the user to define new Syslog facility codes. One of the commonly noticed limitations of Syslog protocol has always been that the "Facility" codes (which define the data source) are limited to 24 predefined codes. The CorreLog program removes this restriction, permitting users to define their own facilities, such as "Applications", and "Devmsgs", so that data can be better categorized and managed. This important extension to the Syslog protocol opens important new vistas in the practical use of Syslog messages not otherwise available using the standard Syslog specification.

- **Ability To Override Device, Severity and Facility Info.** CorreLog allows the user to assign device names, severities, and facility information to messages based upon message content. One of the commonly noticed limitations of Syslog protocol has always been that, since messages are unsolicited, the user is stuck with whatever message, severity, or facility was originally specified by the message sender. Although this is not necessarily a problem, in some cases the severities or facilities within a message may be nonsensical. The CorreLog system recognizes this existing limitation and implements a sophisticated “override” technique, which allows users to override the facility, severity, or device name in any message.
- **Web Based Configuration, Viewing, and Reporting** The CorreLog Server system is entirely web-based. All activities, including the establishment of logins and permissions, are completely achieved without a native console. This means that an administrator does not ordinarily need access to the CorreLog Server platform, except in rare instances to startup or shutdown the process. The location of the CorreLog Server can be strategically placed in a Network Operations Center (NOC) or secure cabinet, which has important implications for security.
- **Auto-Learn Capability.** The CorreLog Server system includes an "Auto-Learn" function, which monitors message reception, and automatically adjusts thresholds and parameters based upon message rates. This allows the program to learn about its environment, and make suitable adjustments to increase the pertinence of cataloging and correlation functions. This powerful and unique function minimizes the amount of tuning needed by operators, and allows the program to run in an "unattended" mode of operation.
- **Distributed Management Capability.** The CorreLog Server can operate as either a single point of data collection, or as a component in a larger management strategy. Each copy of the CorreLog Server can serve as a "collection agent" within your enterprise, feeding this information to a higher-level collection agent. Additionally, each CorreLog provides components that permit the user to search and obtain status information from other CorreLog Servers, permitting multi-tiered implementations that can handle millions of devices and many Terabytes worth of data.

The CorreLog Fast Start

The remainder of this manual discusses the CorreLog Server in detail. For those users wishing a quick start, the following information will get CorreLog up and running as quickly as possible, permitting you to immediately begin using the program.

- After installing CorreLog, you can access it via a web browser at port 80 or port 88, or port 888 or port 8888. The particular port number is displayed and configured in the installation dialog, and is by default port 80, but may also default to one of the other port numbers if port 80 is already being used by some other program. You can typically access CorreLog via a desktop shortcut on the machine where CorreLog is installed.
- The default login to CorreLog is username “admin”, password “admin”. After logging into the system, you can click on the tabs, at the top of the browser display, selecting first “System” and then “Logins”. This will permit you to view and modify the login username or password, or add other user logins.
- To receive messages, you must redirect the Syslog servers of your Unix boxes, your Cisco routers, and / or install the Win32 Syslog Utilities, found in the “system” directory of your installation. In particular, you should install the CO-sysmsg.exe service on each Windows box you plan to manage, as discussed in Section 2 of this manual. This will allow CorreLog to manage your Windows event logs.
- To view the message information, click the “Messages” tab at the top of the display. From that location you can search for messages, view the catalog information, filter and override messages, or set parameters governing message collection.
- To search for messages, use the “Messages > Search” tab for quick searches by keyword. For more rigorous “Forensic” searches, use the “Reports > Query” tab, which permits searching by complex match pattern. (Links to these two tabs are provided at the top right of the web display.)
- To correlate messages, click the “Correlation” tab at the top of the display. The first tab “Threads” forms the basis of most correlation functions, and allows the user to inspect and create classes of messages. The user can click down into any thread and click the “Analyze” button to view the devices, users, facilities, and severities associated with each message class.

Future sections will describe in detail the various other features, adaptations, customizations, and applications associated with the CorreLog Server. The reader is encouraged to experiment with the system.

In particular, almost all of the information required to understand the essentials of the CorreLog Server has now been explained, you can begin exploring your enterprise message information right now.

Section 2: CorreLog Installation

A detailed description of CorreLog installation can be found in the "CorreLog Quick Installation Guide", available from the CorreLog website, and included in the standard CorreLog distribution package. Additionally, the CorreLog Server is a part of a stand-alone web framework, the installation of which is described in within the "CorreLog Sigma Web Framework" Users Manual, Section 2.

What appears in this manual is an abbreviated (but complete) description of CorreLog installation, useful for quickly installing the default version of the program with no variations.

The CorreLog Server program is usually delivered as a self-extracting WinZip file, and contains install and uninstall programs residing in the "system" directory of the CorreLog root directory. The install program normally starts after files are extracted.

To uninstall the system, the operator access the Windows "Add / Remove" programs application, within the Windows "Control Panel", and clicks on the "CorreLog Framework" entry.

CorreLog is specifically designed not to scatter DLL or other files into system directories. All files within the CorreLog directory reside in the CorreLog root directory, by default the directory C:\CorreLog. Therefore, if the user stops the CorreLog Framework service, the entire CorreLog directory can be simply dragged and dropped into the Windows "Recycle Bin" and this will effectively discard the entire installation.

CorreLog System Installation Requirements

The CorreLog system is minimally invasive, and can be installed on a variety of platforms and operating systems. An "Administrator" login is required to install the software. The system is designed to run on a assortment of different hardware and software configurations, including full support for all Windows 20XX platforms.

Specific system requirements of the CorreLog Server are described below.

- **Disk Space.** The CorreLog Framework, by itself, has a small footprint of less than 50 Mbytes. However, the program can become quite disk space intensive, depending upon the amount of data that is logged. It is typical to require one Gigabyte of free disk space, or even more if intensive logging is anticipated.
- **CPU Requirements.** The CorreLog Framework will normally require a dedicated server, depending upon how many messages are to be logged. To high message throughput will require a high performance CPU. If the load on CorreLog is light (that is, a few dozen messages per second) then the CPU requirements are less stringent, and the CorreLog system can co-exist with other server applications.
- **Network Connectivity.** The CorreLog Framework is a web server and web based application. A static IP address is required. It cannot be installed on a platform that does not have TCP/IP connectivity. To achieve high throughput, no other network applications should co-exist with the CorreLog Server. Additionally a high performance network interface card (NIC) should be employed.
- **Service Ports.** CorreLog requires control of a single TCP service port (normally port 80, but possibly any other port selected by the user.) Additionally, Correlog controls UDP ports 514, which is the standard service port for listening to Syslog messages on the network.

To insure proper installation of the program, the user should close all windows, and temporarily disable any port blocking or Virus Scan software on the system. Reboot, after installation, is not required, but is recommended.

Detailed requirements can be found in the "Quick Installation Guide", including a list of all required and optional ports, and specific details associated with operating systems and hardware specifications. (The "Quick Installation Guide" is available from the "Home" screen of the CorreLog Server web interface.)

Basic CorreLog System Installation Steps

1. Login to the target platform with Administrative permissions.
2. If the platform is running port blocking or Virus Scan software, this software should be disabled during the installation. Additionally, the HTTP web service port (nominally TCP port 80) should be opened, as well as the Syslog port number (UDP port 514).
3. Execute the self-extracting WinZip file, and extract files to the target directory, by default the directory C:\CorreLog, but possibly some other location of your selection.
4. When the self-extracting WinZip file completes, the automatic installation procedure starts. The initial screen of the installation program is shown below.



5. Follow the prompts of the automatic installation procedure. You can typically use all defaults during the installation.
6. When the installation dialog finishes, the HTTP server and "CO-Syslog.exe" program will be running on the platform, awaiting messages. You do not have to reboot the server platform.

7. Login to the CorreLog web interface by clicking the shortcut on the desktop. The default user name is “admin”, password “admin”. Optionally, click over to the “System > Login” screen, and either modify the existing “admin” password, or add a new login and delete the “admin” password from the system.

No other steps are needed to install the program.

Installation Checkout

Once the system is running, the user can test the installation using the “sendlog.exe” program, which resides in the “CorreLog/system” directory. Brief help on how this program is used can be acquired by running the program at a Windows command prompt, and typing the command “sendlog”, which will show the syntax of the command.

Send an initial Syslog message to verify the CorreLog Server is listening for messages. At a command prompt on the CorreLog Windows platform, type:

```
sendlog.exe localhost "First Test Message." 7 1
```

This should cause a message from “localhost” to appear in the “Messages” section of the web interface. The severity of the test message will be debug(7), and the facility of the message will be user(1).

If the message is not received, try using the “netstat -a -p udp” command to verify that a process is listening to port 514, which is the Syslog port on the system. If the process is not running, check to see if the CO-syslog.exe process appears in the Windows Task Manager. If that process is not running, try starting the process in the “system” directory, in a command prompt, to see what error messages are displayed.

Finally, make sure that any port blocking or Virus Protection program is not interfering with the proper operation of the CorreLog Web Framework server, or with the CO-syslog.exe program. This is a common problem, but easily fixed by accessing the “exclusions” section of your protection software, and permitting access to these two ports.

Configuring Syslog On Network Computers

Once CorreLog Server is running, all that is required is to point the Syslog agents of various host computers to the location of the CorreLog platform. The precise steps needed to do this are platform dependent, and the user should consult the operational manuals of each managed platform for details. The general guidelines for accomplishing this activity are provided as follows:

- **Configuring Syslog On Win32 Platforms.** Although Windows does not have native Syslog capability, the CorreLog contains components that can be used to convert event messages to Syslog messages. On the "Home" screen of your CorreLog installation, you can download the CorreLog WTS (Windows Tool Set) and follow instructions to install Syslog capability on any Windows 200X, XP, or Vista system. You should install this program on each Windows platform you are planning to manage. (See notes below.)
- **Configuring Syslog On Unix Machines.** Each Unix platform has a native Syslog program, which can be configured to forward messages to the CorreLog platform. This is typically accomplished in the "/etc/syslog.conf" file of the Unix platform. The administrator logs into the Unix platform, and forwards all messages to the Windows platform running the Syslog program. This is typically accomplished by adding a directive to the "syslog.conf" file such as `"*. * @hostname"`.
- **Configuring Syslog On Cisco Routers.** Cisco Routers support Syslog protocol, which can be enabled easily by telneting and logging into the Cisco Router. The specific Cisco commands needed to configure the transmission of syslog messages is dependent upon the Cisco operating system and device. Typically, the operator enters "set logging" options needed to specify the syslog receiver IP address, and logging level. The operator should consult the hardware documentation for detailed information.

Additionally, users should look at other various devices and applications of interest, many of which may fully support Syslog protocol.

Using Syslog protocol is quite simple. For example, a great deal of effective management can be performed with simple administrative scripts, using the "sendlog" utility and a Unix crontab job, or the Windows Scheduler program. The ease of Syslog makes it highly agreeable to home grown administrative software. An abundance of information can be found from the web on Java, Perl and PHP scripts that accomplish specific management and security functions.

Installing The CorreLog WTS On Windows Platforms

The CorreLog "Windows Tool Set" is a standard part of the CorreLog server package. It can be quickly installed on any Windows platform to automatically add Syslog capability to that system. The program uses minimal CPU and disk space, and requires very little configuration. No special authorization or license file is required. The WTS may be downloaded directly from the "Home" screen of the CorreLog system, after login to the CorreLog server.

Installation of the Windows Agent and Windows Tool Set is accomplished as follows:

1. Login to the target Windows platform with an Administrative login. This is the target platform for the installation, i.e. where the Windows Tool Set is to be installed.
2. On the target platform, run a web browser, and connect to the CorreLog system running at your site. Login to the CorreLog Server to access the "Home" screen.
3. On the CorreLog Server "Home" screen, click on the "Download Windows Agent And Tool Set Here" link to download the CorreLog WTS software. (If you are using Windows Explorer or a compatible browser, you may click "Open" to execute the program.)
4. After downloading and executing the WTS package on the target Windows platform, the WinZip self-extractor starts. Extract files to the desired location, by default C:\CorreLog.
5. After files are extracted, the installer program automatically starts. The program requires a single argument, which is the location of the main CorreLog server that will receive Syslog messages.
6. Run the install wizard to completion. After the wizard finishes, the CorreLog System Message Service will be started. You do not have to reboot the platform. You should see a "Startup" message logged at the CorreLog server system.

The WTS program may be uninstalled using the standard Add / Remove Programs facility in the Windows Control Panel. Detailed documentation on the CorreLog WTS software, including how to further configure and refine the installation, is provided on the "Home" screen of the CorreLog Server System.

Note that, in addition to the interactive installer (described above) various other types of WTS packages are available: "wt-agent.msi" file provides a standard Windows MSI installer, and the "wt-silent.exe" file provides a silent installer. These are documented in the "Syslog Windows Tool Set Manual", available from the CorreLog "Home" screen.

Installing CorreLog Adapters

In addition to the basic installation, described here, CorreLog implements an "adapter / plug-in" philosophy which allows the user to extend the range of CorreLog to include interfaces to different systems, such as SNMP and Ping

polling, and WMI data. These adapters may also include specific reporting and action features, such as direct interfaces to McAfee EPO or BMC software.

Users obtain adapter software from CorreLog support, and install this software into the basic installation. (A list of CorreLog adapter download packages, and their manuals, is available from CorreLog support on request.) Each adapter adds new screens, features, functions, and processes. Each adapter has its own specific manual that details any additional steps needed to install the system. A typical adapter installation screen appears as follows:



Adapter installation is typically straightforward, but the user should consult the installation section of each adapter manual for the necessary manual steps that may be needed to finish the installation, such as updating the "System > Scheduler" screen, or executing a second installation procedure. The "View Manual" button on the above screen will view the adapter manual. (The adapter manual is also available, after installation, in the "s-doc" directory of the CorreLog root folder.)

CO-trecv.exe TCP Syslog Message Receiver

Correlog includes a program that permits reception of Syslog over TCP, directly supporting Syslog-NG and Rsyslog hosts. This program is a standard component of the system and does not require installation of a plug-in or adapter. However, program does not run by default, and must be manually enabled. This is accomplished easily as follows.

1. Login to the CorreLog Server, and access the "System > Schedule" screen of the CorreLog Server web interface.
2. Click the "AddNew" button on the above screen. On the "AddNew" screen, add the "CO-trecv.exe" as a scheduled program. The user simply adds the program name with no arguments or pathname, and selects the "start" directive for the program via the "Directive" select menu at the left.
3. Optionally, inspect or edit the "CorreLog\system\CO-trecv.cnf file", which assigns the port numbers for the operation as well as other parameters. This file comes with pre-configured defaults that are usually suitable for normal operation, so changes are not usually required by the operator.
4. Stop and restart the "CorreLog Server Framework", and then verify with the Windows Task Manager that the CO-trecv.exe program is correctly running on the system. Any errors associated with the program are located in the "CorreLog\system\CO-trecv.log" file.

After starting the CO-trecv.exe program as described above, CorreLog will accept Syslog messages on UDP 514 (as before) but also on TCP port 514. This expands the role of CorreLog to accept standard TCP syslog messages, framed with a Line Feed (ASCII 0x10) character, as per the operation of both Syslog-NG and Rsyslog.

Note that the CO-trecv.exe program, in addition to listening for standard Syslog messages, also supports the CorreLog "CO-tsend.exe" program, which is the CorreLog tunnel sender program, documented in the "Windows Tool Set User Manual".

Also, note that the CorreLog "SyslogDefender" program (a CorreLog add-on) is available as a method of receiving IPv4, IPv6, TLS, TCP, UDP messages. This program and its documentation is available from CorreLog support. CorreLog "SyslogDefender" should be considered as a more powerful substitute for the standard "CO-trecv.exe" program for those sites requiring rigorous TCP connectivity and encryption.

Finally, note that TCP is not necessarily recommended for every CorreLog installation. A common misconception exists that UDP is less reliable than TCP;

hence some people assume TCP is superior in all cases. In fact, experience shows that TCP connections have their own issues, which affect performance, security, and reliability of the logging process, and TCP is often not preferred as a form of transmission, except in some special cases (such as when handling a syslog feed from a central remote collector.).

Uninstalling The Entire CorreLog Framework

The CorreLog Framework is uninstalled via the “Add / Remove Programs” windows facility (or the “Program Features” facility on Windows “Vista” systems.) The user navigates to this screen (via the Control Panel) and clicks on the “CorreLog Framework” entry to execute the Uninstall program. The user follows the instructions of the dialog to uninstall the CorreLog Framework system.

Note that, unlike some uninstaller programs, the CorreLog Framework uninstaller does not remove any files on the disk. (The uninstaller removes registry keys and services only.) Following the uninstall procedure, the user must physically remove these files, such as by dragging the CorreLog root directory to the Microsoft Windows “Recycle Bin”. This extra step safeguards any accidental removal of data on the system.

Note that CorreLog does not scatter DLL files into any Windows system directories. Hence, it is sufficient to remove the CorreLog root directory after running the uninstall program in order to remove all CorreLog files.

The user can manually execute the CorreLog uninstaller program at a command prompt by changing working directories to the CorreLog\system folder, and then executing “CO-uninst.exe”.

Additional Notes

The user should refer to the “CorreLog Sigma Web Framework” Users Manual, Section 2 for additional important notes, including procedures on how to configure the system to work with IIS, how to change the location of the C:\CorreLog\directory, and how to rerun the CorreLog Framework Installer in order to accomplish specific objectives.

Those notes comprise a major subsection of the CorreLog documentation, and should be reviewed by all CorreLog program administrators. For notes on the basic system parameters, how to change system logins, and how to configure the program scheduler, see Section 3 of that document.

Section 3: CorreLog Basic Usage

The CorreLog Server comes with a suite of web-based screens that allow users to view and search messages, and configure message filters, overrides, and correlation rules. These screens, as well as specific functionality of background processes and agents, are designed to fully support enterprise wide monitoring of security.

The basic steps for installing CorreLog have been presented in the previous section. Additional information needed to deploy CorreLog is available in the "CorreLog Installation Quick Start Guide" (accessed from the "Home Screen" of CorreLog.) The information present in this current section provides a basic tutorial on the steps needed to make CorreLog work for you within your enterprise, including different monitoring strategies that leverage program features and modes of operation, all with the specific intention of meeting your security and compliance requirements.

Specific usage of each screen referenced herein, including all screen fields, is provided in the "CorreLog Screen Reference Manual" companion documentation. (For detailed information associated with screens and controls, refer to that manual.) This current section references those screens, providing instructions on how best to configure and use CorreLog for the management of your enterprise, including information on how to meet the objectives of regulatory compliance requirements.

Different Modes of CorreLog Operation

CorreLog is designed to support various modes of operation, any or all of which may be pertinent to your site. The basic types of CorreLog operation include (but are not limited to) the following typical applications.

- **Unattended Operation.** CorreLog can execute in an unattended manner, gathering and archiving data for post analysis. This mode of operation exists by simply installing CorreLog and directing syslog traffic at the program. If you have installed CorreLog, you are currently running in an unattended mode.
- **Security Monitoring.** CorreLog monitors system security with out-of-the-box correlation rules. These rules are sufficient to monitor the generic security of your enterprise. Operators can periodically check the program to see what security alerts (if any) have been detected by the program.
- **Compliance.** CorreLog has specific elements necessary to meet regulatory compliance objectives. In particular, CorreLog creates actionable tickets that can be periodically reviewed and closed thereby positively demonstrating to auditors and corporate management that CorreLog is proactively integrated with your enterprise operations.
- **Forensics / IT Search.** CorreLog has specific elements needed to support post-analysis, forensics, and IT search functions. You can search for past log messages, interrogating the historical state of your network. This provides

The particular mode of interest depends upon your specific goals. At minimum, having installed CorreLog, you are proactively providing security by the continuous monitoring of log messages that track user activity. Also, you have established a basis for performing forensics, and taken a major step forward with implementing regulator compliance.

Searching and Querying Message Data

One of the main purposes of CorreLog is to locate message data that matches a specific pattern, keyword, or characteristic. This will typically be one of the first things that CorreLog is used for, after installation. CorreLog provides various facilities and screens that allow the user to immediately search for message data.

- **Messages > Search Facility.** CorreLog implements a high-speed indexed search function that is accessed via the "Search" hyperlink at the top of the screen, or by clicking on the "Messages > Search" tab. The "Search" facility allows the user to quickly review large amounts of data to locate messages by indexed keyword. This is the fastest way to search for data

on the system, but requires the user to start the search with a message keyword. Keywords are automatically derived from incoming messages, and include any word used in any message that does not contain special characters or numbers. A list of current keywords is available on the main "Search" screen.

- **Reports > Query Facility.** CorreLog implements a more rigorous (but slower) search facility that is accessed via the "Query" hyperlink at the top of the screen, or by clicking on the "Reports > Query" tab. The "Query" facility meticulously goes through all messages on the system (identified by a selected range) and compares each message to a potentially complex match pattern. The Query facility runs in background, maintains a history of queries. Although the Query facility is not as fast as the Search facility, the Query function permits the user to search for messages that consist of logical combinations of keywords, wildcards, numeric text, and other qualifiers.
- **Correlation > Threads Facility.** CorreLog allows the user to classify and catalog incoming data by a variety of criteria, including match pattern, device group, facility code, severity, etc. These messages are classified as CorreLog receives them, in real time. Within each catalog, the user can search for data (to some limited extent, defined by the "Max Non-Indexed Search" parameter, found on the "Messages > Config > Parm" screen, by default 100,000 messages.) Various predefined catalogs exist, including catalogs by device, username, facility, and severity. CorreLog "Threads" operate in an identical manner to these predefined catalogs, except the classifications are completely user defined.

The above facilities are easy to get started with. The user can click on the "Search" or "Query" hyperlinks, found in the upper right corner of CorreLog displays, to immediately access either the Search or Query functions. The "Search" function is generally faster, but slightly less flexible than the "Query" function (because the Search function uses an indexed search, requiring the user to start the search with a valid keyword).

Reviewing Ticket Information

With no special configuration, CorreLog will begin opening tickets on your system. These tickets consist of actionable data, and represent the highest level of correlation on the system. As the program executes, it learns about message counts, and will adjust thresholds so that tickets become more pertinent over a period of time.

The most minimal level of attending to CorreLog is to close tickets that are periodically opened. This activity serves as a good demonstration of regulatory compliance to auditors, since it proves that security alerts are reviewed.

Additionally, no specific knowledge of CorreLog is actually necessary to review ticket information. A user can be granted "Ticket" type access (via the "System > Logins" screen) that does not grant other access to security data. Hence, reviewing ticket information can be a group activity of various administrators and security officers. A user can process and review tickets as follows:

1. Log into the CorreLog system and click the "Tickets" tab at the top of the screen. This displays a list of tickets assigned to your login ID. (You may also use the "Assigned To:" menu item to view all tickets, including those assigned to other groups or organizational divisions.)
2. For any ticket of interest, click on the "Related Messages" hyperlink to view the messages that caused the ticket to be opened. This can provide an indication of the exact reason for the ticket.
3. If the ticket appears to be the result of a threshold being set too low, click the "Source Alert Definition" hyperlink of the ticket to review the threshold and test interval for the alert that opened the ticket, and make adjustments accordingly.
4. If the ticket appears to be the result of a true security breach, take external action such as to contact the user referenced in the related messages, or the administrator of the system or network related to the ticket.
5. When the ticket has been resolved, click the "Edit" button and optionally provide a brief description of how the ticket was resolved, and close the ticket.

Rather than closing the ticket, a user can elect to assign the ticket to another user, possibly escalating or demoting the ticket severity, and adding notes to the ticket. When a ticket is closed, it is retained in the "Closed" tab of the system.

The CorreLog ticket system is described in detail within Section 7 of the manual, and includes various features such as the ability to connect e-mail (and other) notifications to ticket changes, and to group close tickets on the system.

Defining New Correlation Threads

One of the main purposes of CorreLog is to organize data into "Threads". This is one type of correlation available to users, and is very simple to implement. CorreLog comes predefined with a variety of useful correlation threads that are generic enough to use without modification. Users can (and typically will) create their own correlation threads, pinning these threads to the top for easy reference, and creating alerts for these threads that open tickets on the system.

CorreLog "Threads" are catalogs of messages that are related by some characteristic, such as a common keyword, device name or group, time, facility code, severity code, or some other message characteristic. For example, a user may be specifically interested in data of a particular type generated by a particular device or group. The procedure for creating a thread is as follows:

1. Log into the CorreLog system and click the "Correlation > Threads" tab at the top of the screen. This displays a list of current threads on the system. You may sort these threads, or click on the thread titles to view messages associated with these threads.
2. To add a new thread, click the "Wizard" button in the upper right of the screen. This starts the "Add New Thread" wizard, which guides the user through the process of adding a new thread to the system. (More experienced users may simply click the "AddNew" button to configure a new Thread.)
3. The "Add New Thread" wizard queries you for the "Thread Title", which appears on the top level "Threads" screen and other locations within CorreLog. A descriptive title for the thread should be provided.
4. The "Add New Thread" wizard queries you for basic qualifiers, such as the IP address, message severity, and time range. Each message received by CorreLog will be compared to these values. If a match occurs, the correlation process continues. By default, the default qualifiers match all received messages.
5. The "Add New Thread" wizard queries you for a match expression. To get started, you can specify any keyword, phrase, wildcard, or logical combination thereof. Each message received by CorreLog, which matches the qualifiers of step 4 (if any) will be compared to this match expression. If a match occurs, the message will be added to the thread.
6. Finish the wizard. The thread title specified in step 3 will appear at the top of the list, and will be "pinned". You can edit the thread via the "Edit" button, including unpinning the thread if so desired.
7. As a final step, you can "Regenerate" the catalog, which will populate the thread with data. (By default, if the thread is not regenerated, only new messages will be added to the thread.) The user can click the catalog title hyperlink, and then click "Regenerate Thread" at the bottom of the screen. CorreLog will prompt for confirmation, and then execute a background process that will scan recently received messages and add them to the thread. This process can take several minutes.

If the user regenerates the thread, as described in step 7 above, CorreLog launches a background process that scans all recent messages to some limit specified by the user. All messages that match the thread specifications are entered into the catalog. While the thread is regenerating, the status line indicates the progress of the operation. The user can leave the screen, and later return to see the progress of the operation or see what messages have been entered into the thread catalog.

As an alternative to regenerating the thread, the user can simply check back to see what new messages have been collected for the thread.

A description of the advanced features of threads (such as the ability to specify a trigger qualifier, and search thread data) is provided in section 5, and in other locations within this manual.

Refining and Auditing Correlation Threads

Correlation threads are the first stage of correlation. (Later stages are discussed below, and within Section 4 of this manual.) The out-of-box correlation threads, and any new threads created by the user, should be audited to see if they contain proper data. It is often the case that small changes may be necessary to a thread, or a new thread needs to be derived from an existing thread with small changes applied. The procedure for auditing and refining a thread is as follows:

1. Log into the CorreLog system and click the "Correlation > Threads" tab at the top of the screen. The message counts for each thread appear to the right of the thread counter. (The "History" count indicates the number of messages received since the thread was configured, whereas the "Count" value indicates the number of messages received since CorreLog startup. Both counters will track each other.)
2. If any thread has excessive counts, it is a good candidate for refinement and modification. Click on the Thread Title hyperlink and inspect the messages being assigned to the thread. If a particular message is not appropriate for the thread or not pertinent, make a note of the message content. Pick a unique word in the message that is not appropriate for the thread title
3. Click back to view the main thread display, and click the "Edit" button for the thread. This displays the "Thread Editor" screen.
4. Modify the "Match Expr" value for the thread by appending an additional qualifier. This qualifier will typically be "and not (patt)", appended to the end of the match expression. This will exclude further messages with the specified keyword from the thread.

5. Click the "Save" button to save the change to the thread. As an alternative, the existing thread can be retained unmodified, and a new thread can be created with the change using the "SaveAs" button. (If the SaveAs button is clicked, the user will also need to modify the title of the thread before it can be saved. See notes below.)

The threads that come with CorreLog, while useful, are fairly generic. It is often useful to make more specific threads, retaining the existing thread. For example, the user can edit the "Admin Login Failures" thread, change the title to "Main Server Admin Login Failures", and specify the IP address for the main server. Clicking "SaveAs" will create a new thread that contains the login failures specific to the main server. The original thread will continue to operate exactly as before.

If any changes are made to a Thread, the user should consider executing the "Regenerate Thread" to update the thread contents with the new match pattern. (The "Regenerate Thread" function is described at the end of the previous procedure.)

Creating and Refining Data Groups

CorreLog provides multiple features for grouping data, or creating groups of groups. Several defaults are provided out-of-the-box, but as part of the administrative activities and setup of the system, the operator should create or refine certain group elements to help organize the data. The following groups are particularly important:

- **Address Groups.** The user can create groups of devices by address via the "Correlation > Config > Address Groups" tab. The address groups are given names that can be used in correlation threads, and which appear in the drop down list of the "Messages > Catalogs > Devices" screen. The operator can view the status of device groups via the "View Groups" link on the "Messages > Catalog > Devices" screen.
- **Thread Groups.** The user can create groups of threads via the "Correlation > Config > Thread Groups" tab. These thread groups match thread titles, and help organize the thread list into sections. The thread group names appear in the drop-down list of the "Correlation > Threads" screen. The operator can view the status of thread groups via the "View Groups" link on the "Correlation > Threads" screen.
- **Ticket Groups.** The user can create groups of tickets via the "Tickets > Config > Ticket Groups" tab. Ticket groups represent who can be assigned tickets, and help organize the data into meaningful groups that are subsequently used in ticket notifications and actions. The operator can view the status of ticket groups via the "View Groups" link on the "Tickets > Opened" and "Tickets > Closed" screens.

These groups greatly assist in organizing the data, and should be managed and maintained by the operator and administrator of the system. The user can set their "Initial Group" preferences via the "System > User Preferences" tab (so that screens reflect the data of highest interest to the logged in CorreLog user.)

Defining Correlation Alerts And Tickets

Once a thread is created, the user can optionally add an alert threshold for the thread counter. (The thread counter appears at the right of the thread title on the "Correlation > Threads" screen.)

Alerts compare correlation counters to limits, and then send syslog messages (of the user's selection) back to CorreLog. Alerts can also open tickets on the system, assigning these tickets to specific users and groups. The procedure for creating an alert and ticket is as follows:

1. Log into the CorreLog system and click the "Alerts" tab at the top of the screen. This displays a list of current alerts on the system. You may sort these alerts, or filter the alerts to find keywords. Any alert that is set is depicted as "red". Any alert that is "green" is not currently set. The current threshold and current value for the alert is displayed next to the alert state color.
2. To add a new alert, click the "Wizard" button in the upper right of the screen. This starts the "Add New Alert" wizard, which guides the user through the process of adding a new alert to the system. (More experienced users may simply click the "AddNew" button to configure a new Alert.)
3. The "Add New Alert" wizard queries you for the "Counter Name", which is listed as a drop down menu. You can select any thread, as well as various other counters on the system. When a new thread is added to the system, it appears in this drop down list.
4. The "Add New Alert" wizard queries you for the compare function, threshold and test interval. You can click on the "View Counter Threshold Hints" hyperlink to see recommendations (and then click back to return to the wizard.) Or, in the absence of any data or other requirement, you can accept the default value of 3 per 60 seconds, and the "Auto-Learn" function will make adjustments accordingly.
5. The "Add New Alert" wizard queries you for a facility, severity, and who to assign the ticket to. Optionally, adjust the severity of the alert. In the absence of any other requirement, use the other defaults for the screen, i.e. assign the ticket to yourself.

6. The "Add New Alert" wizard queries you for an alert message. This is the message that will be sent to the CorreLog server, and is also the text that will appear in the ticket message. The wizard always suggests an appropriate message. However, you may further tailor this message by adding or deleting text. (You can also make these and other adjustments after the alert is created.)
7. Finish the wizard. The alert will appear at the top of the alert list, and will be "pinned". You can edit the alert via the "Edit" button, including unpinning the alert.

Once the alert is created, the user can test the alert and verify it correctly opens a ticket. This can be accomplished using the "More" menu in the upper right corner of the display. Select "Send Msg" from the menu. (This tool is also available at the top of the "Messages > Search" screen.)

Using the "Send Syslog Message" tool, furnish a test message that will match the thread (or other system counter) specified in step 3 above. Note that you can specify the source device for the message via the "Additional Options" of the "Send Syslog Message" screen. Send the message until the threshold specified in step 4 is achieved. The alert will turn red, and a ticket will be opened, which will appear in the "Tickets > Open" tab of the system.

Configuring SMTP Notifications

When a ticket is opened, closed, or modified, it can run an action program such as sending e-mail. Configuration of ticket actions may or may not be an essential activity of your security monitoring. Most sites will send e-mail to an administrator when certain tickets, or all tickets, are opened. Prior to configuring any ticket actions, the administrator must first configure the CorreLog SMTP server settings as follows:

1. Log into the CorreLog system and click the "System > SMTP" tab at the top of the screen. This displays the current SMTP settings of the CorreLog server. (Initially, the SMTP server is not configured.)
2. Click the "Edit" button to edit SMTP parameters. This displays the SMTP Parameter Editor screen.
3. On the editor screen, provide the various parameters associated with the SMTP server. These parameters should be familiar to most administrators and Microsoft Outlook users. You must specify the SMTP server as an IP address, and specify a valid login to the SMTP server. This login will be used for sending all e-mail messages, and must include a valid "Mail From" address.

4. Click "Save" to save the SMTP settings. This will redisplay the screen shown in step 1 above.

To test the SMTP server settings, click on the "Send Test Message" button at the bottom of the SMTP parameters screen. This will send a short message to the "Mail From" user, and will display a transcript of the SMTP server exchange. The transcript should indicate that the message was successfully sent, or should indicate the reason for failure (if any) such as improper authorization due to an invalid password or account name.

Configuring Ticket E-Mail And Notifications

Once the SMTP parameters are configured, the administrator can subsequently configure one or more e-mail actions for tickets. These actions will send e-mail messages to specified user, possibly qualifying the e-mail based upon ticket content, severity, time of day, or ticket assignee. Configuring an e-mail ticket action is accomplished as follows:

1. Log into the CorreLog system and click the "Tickets" tab at the top of the screen. This displays a list of current tickets assigned to you (or to "All" users if you are the administrator.)
2. Click the "Tickets > Actions" screen to display a list of ticket actions on the system. Initially, no actions are configured.
3. To add a new ticket action, click the "Wizard" button in the upper right of the screen. This starts the "Add New Ticket Action" wizard, which guides the user through the process of adding a new ticket action to the system. (More experienced users may simply click the "AddNew" button to configure a new ticket action.)
4. The "Add New Ticket Action" wizard queries for the type of action to add. Select the "SENDMAIL" ticket action, which will send e-mail to a user.
5. The "Add New Ticket Action" wizard queries for a ticket user and state. By default, an e-mail message will be sent for any ticket assigned to any user when the ticket is first opened. You may select a more specific user or different state.
6. The "Add New Ticket Action" wizard queries you for additional match qualifiers, including a possible keyword that must appear in each ticket. In the absence of any other requirements, you can use the defaults.

7. The final screen of the "Add New Ticket Action" wizard queries you for the e-mail address to send the ticket information to. Supply an valid e-mail address and finish the Wizard.

When the action is configured, the user can test the action by simply opening a ticket manually. The user can click the "AddNew" button on the "Tickets > Open" screen, enter the parameters of the ticket (including the appropriate severity and match keywords, if needed) and then click "SaveNew" to save the ticket. This will cause any action programs that match the content of the ticket to be executed. The user can click the "View Debug Log File" for any error indication that caused the action not to run.

Configuring E-Mail Reports

CorreLog provides a substantial reporting facility, with many features and functions suitable for use by analysts and auditors. A complete discussion of the reporting facility is provided in Section 8.

One of the easiest and most useful types of reports to create is an "E-Mail" report. This type of report mails users a short list of messages, of the operator's selection, each day at midnight. In addition to providing an easy way of reviewing pertinent activity, this also demonstrates to auditors that CorreLog is well integrated into the enterprise.

At minimum, the "CorreLog Internal Events" thread should be mailed to the CorreLog administrator. This report lists changes to CorreLog by users, records significant internal events and anomalies, and verifies that CorreLog is running properly. Configuring this particular E-Mail report is accomplished as follows:

1. Log into the CorreLog system and click the "Reports > E-Mail" tab at the top of the screen. This displays a list of all configured E-Mail reports on the system. (Initially there will be no E-Mail reports configured.)
2. To add a new E-Mail report, click the "AddNew" button. This displays the "Add New E-Mail Report" dialog.
3. On the "Add New E-Mail Report" dialog, select "CorreLog Internal Events" as the "Report Data Source". (The user can select any configured thread as the data source via the "Report Data Source" drop-down menu.)
4. On the "Add New E-Mail Report" dialog, provide an e-mail subject, such as "CorreLog Internal Events".
5. On the "Add New E-Mail Report" dialog, provide a valid e-mail address of the user that will receive the daily report.

6. The user can accept the remaining defaults of the screen to generate a report of under 500 messages each day, reporting on the previous day's messages. Click "Save" to save the settings. The new report will appear on the top-level "Reports > E-Mail" screen.

To test the report, the user can click the "Run Report" button at the upper right of the screen. This causes the report to be immediately generated and e-mailed to all specified users. (SMTP server settings should first be configured for e-mail reports to work, as discussed in a previous section.) Additionally, clicking on the E-Mail subject line on the top-level screen will display the report contents, or any errors that occurred during the generation of the e-mail report.

Reducing Correlation And Message Load

CorreLog can continuously accept 2000 or more messages per second, and handle much larger bursts of messages for a short period of time. The exact amount of messages and events per second (EPS) depends on multiple factors such as the number of correlation threads and CPU power of the host platform.

However, it is often the case that only a small fraction of these messages are actually pertinent or necessary for security compliance. In particular, redirecting all firewall data at CorreLog can quickly make the program difficult to use (because so little data is security event data, pertinent for security management)

CorreLog provides a unique method for reducing the message and correlation load, especially applicable to firewall data, but also important for many other types of devices and messages. Using the CorreLog "Filter" facility, the user can arbitrarily redirect messages away from the main message stream into "Auxiliary" files.

Note that CorreLog filters do not necessarily delete data, which may be important to support forensics. Rather, the filters move the data away from the correlation engine and indexed search engine into auxiliary files. The data in this file can be reported on, searched, and archived like any other data.

If the "Messages > Search" screen is filled with non-pertinent messages of a particular type (such as firewall data, or any other message that repeats in a similar fashion for many pages of data) the message load can be reduced as follows:

1. Log into the CorreLog system and click the "Messages > Search" tab at the top of the screen. This displays a list of incoming messages. The most recent messages are listed first.
2. On the "Messages > Search" screen, inspect the list of messages to see if there are many messages with common keywords that look like

candidates for redirection to the auxiliary files. You can use the "Search" function to find the common keywords and data characteristics. Note that it is quite common that one device will be overwhelming the event log with data. Make a note of the device, message severity, and keyword associated with the most ubiquitous and non-pertinent messages.

3. Click the "Messages > Config > Filters" tab to display the list of filters on the system. (Initially no filters are configured.)
4. Click "AddNew" to add a new filter to the system.
5. On the "Add New Filter" screen, enter the match keyword, device address, and message severity noted in Step 2 above. The data should be specific enough to filter only non-pertinent messages. For example, it may not be sufficient to filter only by keyword (since a remote chance may exist that the keyword appears in interesting messages that should be left unfiltered.)
6. On the "Add New Filter" screen, select the "Aux-1" filter file as the "Filter Output" file. (Note that the default "Main" output file should be selected only if the message is to be entirely discarded. The report facility and archive functions work only on "Aux-1" through "Aux-8" files.)
7. Click "Save" to save the filter setting. The new filter appears on the top-level "Filters" screen.
8. Repeat steps 2 through 7 above, as needed to reduce the number of non-pertinent messages. CorreLog can accept up to 1000 different filters. Each filter can redirect data to the same file, or different files. For example, there may be various filters redirecting to the "Aux-1" file.
9. When all filters have been configured, click on the "Messages > Aux" tab to display the current Auxiliary files. The auxiliary files will appear in the list of files, along with the size and message count of each file.
10. Optionally, click the "Advanced" button at the top of the "Aux" display, and provide a descriptive title for each "Aux" file, such as "Firewall Data", or "Miscellaneous Messages". This optional step can clarify the purpose and intent of the Aux file data filters.

Once an Aux file has been defined, the "Size" and "Count" fields on the "Messages > Aux" screen will begin to increment, showing that messages are being redirected to this file. The user can click on the "Aux Filter File" hyperlink to view recent messages or to search messages. The user can click the "Reports > Query" function to run more complex queries across one or more Aux files.

If the "Aux" file is to be retained in the archive, the administrator should additionally navigate to the "Messages > Config > ParmS" screen, and set the "Archive Filter Data" setting to be "Yes". This causes the auxiliary data to be compressed and placed into the archive before it is deleted each night.

The CorreLog "Advisory" Facility

CorreLog includes a built-in "Advisory" function, which performs hourly checks on a variety of system performance metrics, including disk space utilization, process execution, thread counts, message loading, and other checks. If any anomaly or warning indication is discovered during these checks, a special "Advisory" icon and link appears at the top of the screen, and the operator can click down to view the advisories.

No special configuration or setup is required to enable the "Advisory" function. A complete list of advisories can be viewed, enabled, and disabled via the User Preference screen (accessed via the "System > Prefs" tab) by clicking on the "Modify Advisories" button towards the bottom of that screen. By default, all advisories are enabled for all users. An operator can acknowledge advisories, and can disable specific advisories in their logon profile.

Advisories are distinguished from tickets and other events, in that advisories apply to CorreLog system performance and health (as opposed to tickets, which generally pertain to the security and health of monitored systems.) Various built-in health functions exist, and the administrator can post an advisory via the "More > Post Advisory" screen, which appears in the advisories of all users of the system.

Generally, any advisory of the system should be investigated and addressed as stated in the advisory text. Although an operator is free to disable any advisory, this should be done cautiously in order to avoid masking a potentially serious system health issue or security risk. The operator can see a list of "Recent Advisories" via a link on the main advisory screen, and within his or her user preferences. Users should occasionally audit this list of recent advisories to see if there are any advisories of interest that the user may have missed.

Note that advisories are generated at the beginning of each hour, but acknowledging any advisory suppresses further updates for the user for at least one hour. Therefore, it may be near two hours before an advisory condition is detected. For example, if a user clears his or her advisories at 2:30pm, the next advisory will be received at 4:00pm.

Local System Advisories

The Advisory function includes a "Local Advisory" capability, which permits a system administrator to post an advisory message to all users of the system. The

administrator accesses this function via the "More > Post Advisory" screen (in the upper right corner of the CorreLog display.)

Each local advisory can contain up to 250 characters of arbitrary text, and each advisory contains an "expiration time" in hours. (After the specified expiration time, if a user has not yet acknowledged their advisories, the advisory is eliminated on the system.)

The "Local Advisory" function furnishes the administrator with the capability to issue advisories that affect workflow, such as warning indications that may exist in the enterprise associated with emergent threats. Unlike advisories generated by the system, which are generated each hour, local advisories are immediately displayed for all users. Otherwise, local advisories operate in a manner similar to the built-in advisories.

Note that only one local advisory exists at a time. Issuing a local advisory removes any previous local advisory (but does not affect other advisories that may exist on the system.) Therefore, if an administrator wishes to eliminate an advisory from the system, the administrator simply posts a message such as "Ignore", with an expiration time of one hour. Any local advisory that previously existed will then be overwritten by this new advisory.

Auditing Recent Advisories

If a user has disabled a particular advisory, the complete list of advisories is still available to the user via the "Audit All Recent Advisories" link. This link appears at the top of the System Advisory Screen, and can also be accessed from the "User Preferences" screen.

The "Audit All Recent Advisories" link allows a user to see advisories that may be disabled in their preferences. In particular, this affords a way to see if the user is missing important advisories, allowing the user to re-enable these advisories in their preferences. Advisories are listed in chronological order of occurrence, and contain the number of times that the advisory occurred, along with any information specific to the advisory.

All advisories that are older than the "Keep Data" setting (by default 30 days) are dropped from the audit list.

Section Summary And Additional Notes

1. CorreLog employs a tabbed navigation system, where certain tabs are nested within other tabs. Users can traverse the navigation tabs to reach the appropriate screen of the system.

2. The "Search" screen allows the user to view recent (and also historical) message data. The messages are displayed in reverse chronological order, so the most recent event is displayed at the top of the screen.
3. The "Correlation" tab provides access to special correlation features of the CorreLog Server; in particular user defined "Correlation Threads". These comprise a major part of the CorreLog Server, and are fairly easy to use
4. The "Alerts" component monitors catalog counters, and can send a message back to CorreLog stream for further correlation, and can open tickets when incidents occur.
5. The "Tickets" component provides the highest level of correlation on the system, and creates actionable tickets based upon Alerts. Each ticket is associated with a particular alert.
6. The "Tickets" component permits the user to open and close tickets. Tickets can run actions, such as sending e-mail.
7. The CorreLog E-Mail interface, accessed via the "System" tab, allows the administrator to configure the SMTP parameters used by e-mail notifications and other sections of the system.
8. The operator should monitor the system load, and reduce the amount of information being correlated. In particular, firewalls and other chatty managed devices can have their messages redirected to an auxiliary file to reduce the correlation load, using the "Filters" section of the program.
9. CorreLog includes an "Advisory" interface that performs hourly checks of system health, and which create an "Advisory" link at the top of the display when a problem exists. Users can acknowledge and disable specific advisories in their profile.

Section 4: CorreLog Applications

The previous section presented the basic CorreLog screens. This section discusses how to use those screens, elaborating on three principle areas of CorreLog functionality, as follows:

- **Data Collection and Archiving.** CorreLog stores message data in organized modules that permit processing by scripts and third party software. This data is held online for up to 500 days. Additionally, the CorreLog system can archive data in compressed format for more than 10 years (5000 days).
- **Interactive Data Browsing.** CorreLog provides a suite of web-based screens, including an indexed search engine that permits users to browse data. In particular, the search engine makes it simple to locate information by device, facility, severity, or message keywords.
- **Correlation and Automated Actions.** CorreLog provides methods to take action when specific messages occur, such as relaying specific messages to other Syslog servers, network managers via SNMP, or running specialized programs to update databases and perform other notifications. The CorreLog Server provides out-of-box utilities, as well as an easy API to permit construction of other automation tools

The above topics are the focus of this section, which will be useful to end-users and operators, as well as system managers and developers. The narrative provides an expansive and detailed overview, and provides pointers to other sections and appendices as reference.

CorreLog Services And Processes

CorreLog employs various processes that collect information, all of which run as persistent background services on the host platform. These programs cooperate to gather Syslog data, organizing this data into catalogs.

All executable programs reside in the CorreLog/system directory, and are started by the CO-svc.exe program on node startup, or when the CorreLog service is started manually. For proper operation of the CorreLog Server, the following processes should continuously run, and be visible in the task manager.

- **CO-Syslog.exe.** This program listens to UDP port 514, the standard Syslog port, and is responsible for receiving and writing messages to the "CorreLog/logs" directory. The process is designed for speed, and can process 40 incoming messages per second, or a burst of 1000 messages in a single second. The only other function this process performs (other than getting data from the UDP port and writing it to the disk) is the management of filters and overrides configured by the operator.
- **CO-Gendex.exe.** This program indexes the data received by the CO-Syslog.exe program, to support the high-speed search engine. The program maintains the files in the "CorreLog/logs/dex" directory.
- **CO-Devlog.exe.** This program manages various screens in CorreLog, including the "Devices" screen, "Facilities" screen, and "Severities" screen. The program does not listen to any port, but continuously reads the latest syslog data and writes this to the "catalogs" directory. The program complements the CO-catlog.exe program, discussed below.
- **CO-Catlog.exe.** This program "catalogs" the data (as discussed further below.) It does not listen to any port, but continuously reads the latest Syslog data and writes this data to the "catalogs" directory, maintaining index files to this data and updating CorreLog Server counters. The process also is responsible for executing any actions on the data, discussed later in this section.
- **CO-Action.exe.** This program manages the "./queue" directory of CorreLog and is responsible for running action programs (and potentially other programs) required by the system. The process should be running on all platforms, unless otherwise specified by CorreLog support.
- **CO-Svc.exe.** This program starts the other processes on the system (including the Apache server, if so configured) and can be used to launch other periodic and transient programs hourly, daily, weekly, or monthly. The CO-svc.exe program is the actual interface to the Windows Service Manager.

Other processes, such as the CO-Maint.exe program, are also launched periodically, and are responsible for cleaning log files and catalogs, archiving data, generating reports, and publishing data via RSS. In particular, the CO-Maint.exe is run nightly, just after midnight, and is responsible for cleaning up log files and catalogs.

The above programs form the core CorreLog Server background processes. They are started and stopped by the Windows Service Manager. In particular, the system can be started via the “net start CorreLog” command. Additional processes (such as the GenDex.exe search engine index creation program) are occasionally executed by the CO-svc.exe program, as needed.

Syslog Process Overview

Without any other facility or service required, the CO-Syslog.exe process can be used to gather Syslog data for use by third-party software, user scripts, or simply to archive Syslog data for historical analysis and data aggregation purposes.

The CO-Syslog.exe program writes Syslog messages to the “CorreLog/logs” directory in clear text and easily parsed format. The log files are specifically kept uncompressed and easily viewable to support user scripting. (The “DEX” index files, and the catalog data is generally not available for scripting, however is in an open format agreeable to more sophisticated programming and scripting.)

Users familiar with standard Syslog protocol will have no problem immediately recognizing the format of this data. (Users that are unfamiliar with Syslog protocol should refer to the Appendix section for a description of this protocol.) Experienced users will note that the format of the messages, within the log files, is slightly atypical. Specifically, the timestamp of the message is represented as an ISO time value, which is supported by most databases, and which is easily sorted and parsed. Additionally, the facility codes and severities are written in clear text rather than numerically. This makes parsing, scripting, and searching data with utilities such as “grep” or “find” much easier for a programmer or administrator.

Syslog File Naming Conventions

Within the “logs” directory, Syslog records are written in chronological order, continuously appended with information. Each file in the logs directory is given the name “yyyy-mm-dd.log”, representing the date that the Syslog information was collected. At midnight, the file is closed, and not touched again by the Syslog process. A new file is started for the particular date. This file naming convention makes it very easy to sort the data, or drop old data from the system (such as via the CO-Maint.exe program, discussed later.)

Within the Syslog files are records, one message per line, consisting of the following information.

- **Date.** The first space-delimited field of the message is the date that the message was received, and is in yyyy/mm/dd format. This corresponds to the date of the file name, and is included to permit easy processing by third-part scripts.
- **Time.** The second space-delimited field of the message is the time that the message was received, and is in hh:mm:ss twenty-four hour format. Note that this is the time at the CorreLog Server that the message was received, and exists in addition to any timestamp within the message content. The time value reflects the time zone of the host platform, which many sites like to set to GMT, but which can easily be a local time as well.
- **IP Address.** The third space-delimited field of the message is the IP address of the device that sent the message. It is the address detected by the CorreLog Server, and exists in addition to any address within the message content. This IP address value can be overridden by the “Messages > Configure > Overrides > Device” screen, replaced with an arbitrary IP address based upon certain patterns that match the message.
- **Facility.** The fourth space-delimited field of the message is the textual “facility” contained in the message. Each Syslog message requires a facility to be specified. (If for some reason a message is received without a facility code, the message is still logged, and a default facility of “internal” is specified.) The list of text values corresponding to the facility codes is found in Appendix A of this document. This facility can be overridden by the “Messages > Configure > Overrides > Facility” screen, replaced with an arbitrary text string based upon certain patterns that match the message.
- **Severity.** The fifth space-delimited field of the message is the textual “severity” contained in the message. Each Syslog message requires a severity to be specified. If no severity exists, the severity of “debug” is assumed. There are eight severity codes: “debug”, “info”, “notice”, “warning”, “error”, “critical”, “alert” and “emergency”. The severities are defined by the Syslog standard discussed in Appendix A of this document. The severity can be overridden by the “Messages > Configure > Overrides > Severity” screen, replaced with some other severity based upon certain patterns that the message.
- **Message Content.** The remainder of the message, to the end-of-line, is the arbitrary message content. The Syslog standard specifies a maximum of 1024 characters, but the CorreLog Server will actually log up to 5000 characters per message before truncating the message content. Within

the message may be any clear text ASCII characters, as well as non-ASCII control characters. The content is defined by the developer / originator of the message.

The first five fields are space delimited. Following these fields (to the end of line) is the Syslog message content, which is completely arbitrary in nature.

Relocating The Syslog Log Directory

Note that the “logs” directory can consume large amounts of disk space. For this reason, it is common to relocate the “logs” directory to a separate dedicated disk. For best results, the disk should not be a network server (because this can dramatically slow down the CO-syslog.exe process.)

The user specifies the location of the new Syslog data directory on the “Syslog > Configure > Parms” screen of the web interface. For example, the administrator can specify an auxiliary disk, such as the D: drive. Note that (as annotated on the Parms screen) this value does not take effect until the user restarts the Syslog process, or the node is rebooted.

Changing The Syslog UDP Port Configuration

It will generally never be necessary to change the UDP port or Syslog configuration. This is because all the information of interest is occurring at this port, and most Syslog message generators do not give users the option of changing the port.

If it becomes absolutely necessary to change the port number, a user can hand edit the “CorreLog/system/Syslog.cnf” file, which contains the UDP port (and may contain other directives, not documented here.) The user can stop the CO-Syslog.exe process by stopping the CorreLog service, and then edit the configuration file and restart the service. The file is read only during CO-Syslog.exe program startup.

Enabling Syslog TCP Reception

By default, CorreLog listens only to UDP messages. The user can easily configure reception of Syslog over TCP by enabling the “CO-trecv.exe” program, which is a standard program within CorreLog that does not require installation of any program or plug-in. A discussion of how to enable the “CO-trecv.exe” program using the “System > Schedule” screen is provided in Section 2 of this document.

The CO-trecv.exe program provides two purposes: (1) the program accepts standard TCP messages generated by the Syslog-NG and Rsyslog programs

(found on a variety of platforms, especially Linux platforms), and (2) the program supports the "CO-tsend.exe" program, documented in the "Windows Tool Set User Manual", to allow encrypted tunneling of Syslog and SNMP Trap messages.

The CO-trecv.exe program uses a single configuration file, which is read only once on CorreLog Framework Startup. This file is the "system\CO-trecv.cnf" location, and the file can be edited via "notepad" or any other text editor. The file contains the following directives.

- **ListenPort.** This is the TCP port that the program listens for tunnel messages sent by the CO-tsend.exe and "TUNNEL.exe" programs. The default value is TCP port 51462. The value should generally not be changed without exercising caution, since this is the default port number used by other parts of the CorreLog system.
- **ListenPort-Aux.** This is the TCP port that the program listens for Syslog-NG, Rsyslog, and other Syslog over TCP messages. The default value is TCP port 514, which is commonly (but not consistently) used for TCP Syslog transmissions. It may be necessary for the operator to change this value, depending on the Syslog over TCP implementation being used at a site.
- **TrapDestPort.** This is the location where SNMP traps are relayed to, by default the UDP port 162 on the local machine. This configuration item is required, but used only when the "CO-systrap.exe" adapter is installed at the CorreLog site, and only when using the "CO-tsend.exe" program to collect SNMP traps at a remote location. Otherwise, the value is not used.
- **SyslDestPort.** This is the UDP port used by the CorreLog CO-syslog.exe process, which is the central collector of the CorreLog server. The default value of 514 should not be changes without also changing the contents of the "system\syslog.cnf" file (to change the port at which CorreLog listens for external UDP syslog messages. Except for very special circumstances, the default value of 514 will be correct and appropriate for all implementations.
- **MatchAddress.** This is a match pattern that matches the IP address of incoming messages. The default value of "*" matches any client device. For extra security, this value can limit message reception to a specific device or subnet, useful if there is a small set of CO-trecv.exe clients on the system.
- **ErrorSeverity.** This is the standard syslog severity value that is used should the CO-trecv.exe program encounter an error, such as a bad encryption key, or other anomaly. The default of severity of "error" will be appropriate for most implementations.

- **LogLocal.** This directive is by default "False". If the operator sets this value to "True", then each received message will be logged to the "CO-trecv.log" file (along with other program information.) This is useful for debug, but can generate large files, so the value should normally be "False" to prevent logging of this detailed data.
- **EncryptKey.** This is the default encryption key "seed" value used by the "CO-tsend.exe" program. The value must agree with the configuration found in the "CO-tsend.cnf" file of the remote program. Generally, this value does not need to be changed. The value should only be modified if all "CO-tsend.exe" keys are also modified. (CorreLog uses a pseudo-one-time pad type of algorithm, so the actual encryption key is somewhat irrelevant to the encryption process.)

As noted above, the CorreLog system requires TCP ports 51462 and port 514 to be available for receiving messages. The operator should check to see a firewall or virus protection program does not block these TCP ports.

If any of the above directives of the CO-trecv.cnf" file is modified, the CO-trecv.exe program must be stopped and restarted, such as by cycling the "CorreLog Framework Service" via the Windows Service Manager.

Finally, note that the "CO-trecv.exe" program logs its status and error information to the "CO-trecv.log" file, found in the same directory as the CO-trecv.exe program. This file should be examined if the CO-trecv.exe program fails to start or abnormally exists.

Other Syslog Configuration Files

In addition to the above configuration file, the CO-Syslog.exe program makes use of several other configuration files in the "CorreLog/config" directory. These files are maintained by the CorreLog web interface, and they should not generally be modified or hand-edited except under the specific direction of customer support. Specific files required and used by the CO-Syslog.exe program are as follows:

- **Filt.cnf File.** This file contains the filters, configured at the web interface, which are immediately applied to any incoming data. The file contains space-delimited records, with one filter per line. The fields are: filter time, filter hour span, filter address, filter facility, filter severity, filter keyword, and filter count. The "filter count" is cleared when the CO-syslog.exe program is launched, and indicates the number of times that a filter has been used since startup.

- **Ipadd.cnf File.** This file contains the Address Overrides, configured at the web interface. The Address Overrides are applied to incoming data after filtering. The file contains space-delimited records, with one override per line. The fields are: match address, match facility, match severity, match keyword, new address.
- **Facil.cnf File.** This file contains the facility overrides, configured at the web interface. The facility overrides are applied to incoming data after filtering, and after any Address Overrides. The file contains space-delimited records, with one override per line. The fields are: match address, match facility, match severity, match keyword, new facility.
- **Sever.cnf File.** This file contains the severity overrides, configured at the web interface. The severity overrides are applied to incoming data after filtering, and after any Address Overrides, and after any facility overrides. The file contains space-delimited records, with one override per line. The fields are: match address, match facility, match severity, match keyword, new severity.

Catalog Information

The "CO-Catlog.exe" program is responsible for keeping track of the data, and launching any actions (such as to relay information, or performing special custom actions such as updating a database table with messages.) The data is collected to the "CorreLog/catalogs" directory in an elaborate system of subdirectories and index files required by the CorreLog Web interface. The data in these files is generally not usable by programmers.

The CO-Catlog.exe program requires the config/action.cnf file, which contains the list of actions. The program also makes use of the "./config/thread.cnf" file (if it exists) that stores the user definitions of any custom "Correlation Threads", which are user supplied match patterns. All of this configuration data is editable via the web interface, and the precise format of this file is available on request from customer and technical support.

Archived Log File Information

The "CO-maint.exe" program (which runs nightly) is responsible for maintaining the "archive" directory, which contains archived log files (in Gzip format) for the system. The number of days to keep this information is configured in the "Parms" screen of the system. The user can keep up to ten years (5,000) days worth of information. In addition to archiving log files, the CO-maint.exe program maintains message digests for each log file, useful for forensics.

The CorreLog archiving function, in addition to providing long-term storage of message information, has special functions needed to support auditing and

forensics, such as for HIPAA, PCI, and SOX compliance. More information on the Archive function is provided later in this chapter. Archived files can be searched via the "Reports > Query" screen, documented elsewhere.

Log File Report Information

In addition to archiving data, the "CO-maint.exe" program mentioned above is responsible for the nightly generation of Excel reports. This information can be published to RSS aggregators, or simply stored on the system for historical purposes. The system can keep up to 500 summary reports, generated each month, for a total of more than forty years of data.

The reporting facility is very simple to get started with, but this facility provides a large amount of flexibility for those sites requiring specific reporting functions. Detailed information on the CorreLog reporting facility is provided in Section 7 of this manual.

CorreLog Interactive Usage

Although the CO-Syslog.exe program is quite usable without any other tools, the CorreLog program provides a highly useful interactive interface that will be the target of attention by administrators and users. This interface can easily be navigated by clicking on the tabs at the top of the display. All CorreLog screens related to the Syslog processes (and related to this manual) are located in the "Messages" and "Correlation" tabs at the top of the screen.

Search Messages Screen

The first screen displayed when the user clicks on the "Messages" tab is the "Search Messages" screen, which is one of the most useful screens in the entire suite of CorreLog tools. This screen is the entry point to the high-speed GenDex search engine, which permits searching large numbers of messages by keyword.

It is important to note that this search function distinguishes the CorreLog Server as a true enterprise server, capable of searching gigabytes worth of data within a second or less. Without this capability, operators will be seriously constrained in finding related data items, especially given the volume of Syslog messages typically gathered on an enterprise. Using the "Search Messages" screen, operators can achieve in less than one second a search that might literally take hours using other techniques.

The GenDex Search Engine

At the heart of the "Search" screen is the GenDex search engine, and the DEX files (Data Extraction Files), which are used to index the messages. These files

reside in the “CorreLog/logs/dex” directory, and are maintained by the GenDex.exe program, which is periodically launched by the CO-svc.exe program.

The “Search” screen permits real-time generation of these files without losing search capability. As new messages come into the system, they are searched linearly. Over time, through periodic invocation of the GenDex.exe program, the index files are added to until the files are up-to-date at the end of the day.

If the DEX files become corrupted, they are automatically regenerated and fixed, although this can be CPU intensive if the DEX files (and Syslog files) are quite large. The GenDex program can also be run manually, as discussed in an Appendix to this manual.

Search Results, Start Time, And Page Numbers

For any search initiated by an operator, the entire results set is available to the operator using the “Page” hyperlinks on the screen. Because there may be 1000s of pages, to permit even more rapid searches, the “Search” screen permits the operator to confine a search to a particular start date, and before. This allows an operator to rapidly traverse data through either pages, or setting of the start time.

The “Max” setting controls the particular page size. The default value is Max-50, which shows 50 entries per page. The user can set this value larger (at the expense of rendering time of the browser) or smaller. In particular, the “Max-1” setting permits the user to rapidly assess how many total items exist in the results set.

Search Terms

The “Search” screen accepts keywords and wildcards as follows: the user can search for an IP address with a partial match, however if the keyword is not an IP address, the FIRST keyword must not be partial matches. Following the first keyword, the user can enter other keywords and wildcards, in which case the results set will match the keyword and all subsequent patterns.

Search terms are not case sensitive. They must be more than two characters, and cannot contain numeric values (unless the user is searching for an IP address value.)

These rules are quite similar to other search engines, but contain special features to permit the program to work with message data. The best way to learn the search term algorithm is to try various searches to see what matches are available on the system.

Advanced Search Term Notation

To compose advanced search terms, the user clicks on the “Advanced Search” hyperlink at the top of the “Search Messages” screen, which permits specification of highly qualified search patterns. The search terms composed by this screen follow a simple standard. The rules are as follows:

1. First the search terms are entered, separated by spaces. There can be up to five different search terms. Each term must match completely a word in the result. None of these “Primary keywords” can contain a wildcard.
2. Optionally, the search terms are separated by a “/” token, followed by up to five partial matching terms. These search terms can contain wildcards to span multiple characters. In addition to the primary search term, each partial match must exist in the result.
3. Optionally, the search terms are separated by an “-” token, followed by up to five exclude terms. These exclude terms can contain wildcards to span multiple characters. Any line containing a partial match or wildcard is excluded from the result.
4. Optionally, the search term is appended a “t=” token, followed by the match hour, a “:” sign, and the hour range, so that 1:8 matches any time from 1:00 AM to 8 hours later, at 9:00 AM. The hour ranges from 0 to 24, as does the hour range. If the user specifies this, then any line that does not match the time range is excluded.
5. Optionally, the search term is appended with a “m=” token, followed by an IP address in dot notation, with possible wildcards, such as 127.*.10.*. There must be four dot-delimited fields, each field containing a valid number or asterisk. If the user specifies this, then any line that does not match the specified IP address is excluded.
6. Optionally, the search term is appended with a “f=” token, followed by a formal facility name. The name must be a valid facility. If the user specifies this, then any line that does not match the specified facility is excluded.
7. Optionally, the search term is appended with a “s=” token, followed by a formal severity name. The name must be a valid severity. If the user specifies this, then any line that does not match the specified severity is excluded.

The above rules can be intermixed on a search line to provide a large amount of capability in refining searches of management data. The user can select a keyword, partial matches, excludes, and match the various fields in the message

header (address, facility, severity) as well as confine the search to a particular period of time, such as between 8:00 AM and 5:00 PM.

Search Term Notation Examples

The various search terms provide a good deal of capability, including capability not normally found in web search engines (such as the ability to match “partial” terms, and the ability to specifically match addresses, facilities, severities, and time of day.) Several examples of using advanced search terms, illustrating the above rules, are provided below.

(Note that the “Advanced Search” screen constructs these search terms automatically; hence the description of these rules is mainly of interest to special application developers, and for completeness.)

X Y Z	Match all lines that contain all the full words X, and Y, and Z. Note that XX or YY or ZZ is not matched.
X /Y Z	Match all lines that contain the full word X, and any Y or Z characters in the line.
X - Y Z	Match all the lines that contain the full word X, but not any Y or Z characters in the line.
X /Y -Z	Match all the lines that contain the full word X, any Y characters, but no Z characters.
X t=13:23	Match all the lines that contain the full word X and a timestamp ranging from 1:00 PM to 12:00 PM the next day.
X f=user s=debug	Match all the lines that contain the full word X and the facility user, and the severity debug.
X Y Z /A*C -D*F Q*S a=10.19.*.*	Match all lines that contain all the full words X, and Y, and Z, and match the patterns A*C but not the patterns D*F, or Q*S, and match the address 10.19.*.*

In the above table, the values of X Y Z A C, etc each represent full words, space delimited. The primary keys of X Y Z in the above table must be at least two characters, and cannot contain punctuation marks or wildcards.

Note that all the match patterns require the first token to be a full word match. The user cannot search for a time range without first specifying at least one full match keyword. (This is enforced by the “Advanced Search” screen, but may be confusing to an operator trying to compose the search terms manually.)

Clearing The Search Screen

As a special feature of the Search screen, the operator can temporarily clear the display of events using the “Clear” button. This is a way of “time marking” the display so that only new messages are seen, useful if the operator is awaiting specific events. The operator can still search the data, but clicking on any page link (or on the “Search” tab at the top of the display) will “unclear” the display. The “Clear” function requires the operator to click the “Apply” button in order to refresh the list of messages.

Note that the “Clear” button does not actually acknowledge, clear, delete, or otherwise affect any data logging. It is strictly a way of temporarily clearing the screen of message, to determine current and new activity of the CO-syslog.exe and CO-systrap.exe programs.

Matching Severities

The “Advanced Search” function, and various other screens of the system, permit the user to match a specific severity, such as “debug” or “critical”. However, because the severities have ranked priority values (specifically “debug”, “info”, “notice”, “warning”, “error”, “critical”, “alert”, and “emergency”) unlike facilities or IP addresses, it is possible to match severities based upon a compare function of EQ (equal) NE (not equal) GE (greater than or equal) or LE (less than or equal.)

- **EQ. (Equal To).** This is the default compare function. Match only those messages that have a severity equal to the specified match severity value. For example, the “EQ-notice” setting matches only those message with a “notice” severity.”
- **GE. (Greater Than Or Equal To.)** Match only those messages that have a severity greater than or equal to the specified match severity value. For example, the “GE-critical” setting matches those messages having “critical”, “alert”, or “emergency” severity.
- **LE. (Less Than Or Equal To.)** Match only those messages that have severity less than or equal to the specified match severity value. For example, the “LE-notice” setting matches those messages which have “debug”, “info”, or “notice” severity.
- **NE. (Not Equal To.)** Match only those messages that have a severity OTHER than the specified match severity value. This may be useful to

reject messages that are not pertinent, for example setting a value of NE-warning eliminates those messages that have warning severity.

The above qualifiers permit matches that span a range of severities. In combination with the severities, a total of 32 different combinations are possible (i.e. eight severities X 4 compare functions = thirty two.)

“Severity matching” is provided by a dropdown menu next to the severity value, Note that some screens may not support this function, or may support this function only partially. In this case, the dropdown menu, next to the “Match Severity” pattern, will not be visible.

Managed Device Identification

Each message in the system is associated with a unique device identifier, typically the IP address of the device, but potentially some other keyword such as the unique device name or label for the device. The device identifier for each device is found on the "Messages > Catalogs > Devices" screen.

By default, the Device Identifier is the IPv4 address of the device that sent the message, which appears on the "Messages > Catalogs > Devices" screen of the system, and appears as part of all message screens. This identifier / address is always a hyperlink can be clicked to display the "Device Information Screen", described below, which allows the operator to inspect (or provide) data associated with the device, including optional name, description, associated URL, and settings.

In some cases (such as in a DHCP environment, or when using IPv6 addresses) the user may wish to substitute the IP address for some other value, such as the hostname, using Address Overrides. In this case, the user must first modify the behavior of CorreLog by setting the value of "Require Standard IP Addresses" (on the "Messages > Config > Parms" screen) to be "No" (which disables error checking of IP addresses during message reception.) The user can then override the message address with some other value.

More information on this function is available in this section. (See " Auto-Overriding Device Names" discussion, provided elsewhere in this manual.) This function is likely required for any CorreLog implementation that uses DHCP, and is also useful at sites that wish to eliminate or substitute certain IP address references with more user-friendly labels.

Device Information Screen

Since one of the most useful things to know about a message is the nature of the device that sent the message, the CorreLog provides a “Device Information”

function, which is accessed by clicking on an IP address hyperlink throughout the system. The Device Information screen provides various fields, including the following:

- **IP Address.** This is the IP address of the device, identical to that displayed on the calling screen. In some situations (described below) this can also be a device name, depending upon the setting of the "Require IPv4 Addresses" switch (on the "Messages > Config > Parms" screen)
- **DNS Device Name.** This is the official Domain Name Server name for the device, acquired via a DNS lookup. If the device does not have a domain name, the value is "None". Note that if DNS services are slow, or the device does not have a registered name, the access time of the screen can be degraded. (See DN Helper Program, below.)
- **Ping Response Time.** This is the time for the ping packet to make its round trip between the device and CorreLog Server platform. The time is useful in determining the general response time of the target device and network. The value is available only if the device is accessible via a standard ping.
- **Associated URLs and Extra Tabs.** The administrator can configure miscellaneous URLs, and also extend the tabs associated with the device, so that the operator can conveniently click down on information that may be related to the device. For example, the administrator can configure an optional and arbitrary URL to the device that may reference the web interface of the device, a lower-level CorreLog Server (when operating in a multi-tier operation) or any other URL that is associated with the device.

DN Helper Program

The DNS Device Name is displayed beneath the IP address (i.e. device identifier hyperlink) wherever that value appears on the system. To add a DNS device Name, the user provides a value on the "Device Information Screen" described above.

On large systems, with many devices, the user may spend a lot of time drilling down into devices and providing a DNS name for each device. To facilitate this activity, CorreLog includes the "DNHelper.exe" program in the "CorreLog\system" directory. This program, when executed, automatically fetches the Domain Name for each device, and applies this value to the system.

The "DNHelper.exe" program can be executed at a command prompt, or can be executed at scheduled intervals (via the "System > Scheduler" screen, discussed elsewhere.) The "DNHelper.exe" program requires a single argument as follows:

- **DNHelper.exe –merge.** When executed with the "-merge" option, the DNHelper.exe program will update any IP Address with the DNS name for the device, providing that the device does not already have a name. (If the device already has a configured name, the existing name is not modified.)
- **DNHelper.exe –replace.** When executed with the "-replace" option, the DNHelper.exe program will update all IP addresses with the DNS name for the device, regardless of whether the device has already been assigned a name. This option replaces all DNS names with their current value, as needed.

Note that, depending upon the CorreLog environment and response of the DNS server, it may take several seconds to update each device. Therefore, on systems with many thousands of devices, execution of the DNHelper.exe program may take several hours to execute. If the "DNHelper.exe" program is configured for periodic execution on the "System > Scheduler" screen, the administrator may wish to execute this program no more than once each day at midnight (by setting the "Daily" execution mode for the program.)

SNMP Device Information

Additionally, if the user has installed the "CorreLog/net-snmp" tools, the Device Information screen attempts to perform several SNMP queries of the device. The read community used for this is configured in the "Messages > Configure > Parms" screen (along with the location of the "net-snmp" tools on the system.) The read community must be configured and net-snmp tools installed in order to see the following additional fields on the Device Information screen.

- **SNMP Description.** This is the authoritative description for the device acquired from the SNMP agent. Specifically, it is the value of the "sysDescr.0" SNMP MIB object.
- **SNMP Up Time.** This is the up time of the device acquired from the SNMP agent. (It is actually the uptime of the SNMP agent itself, but can generally be taken to be the uptime of the device that started the agent on boot up.) Specifically, it is the value of the "sysUpTime.0" SNMP MIB object. The field value is displayed in "Timeticks" (100th second intervals) and also in "hh:mm:ss:hh" format, i.e. hours, minutes, seconds, and hundredths of a second.
- **SNMP Contact.** This is the system contact of the device acquired from the SNMP agent "sysContact.0" MIB object value. This value is usually configured by the device administrator, and may be blank or erroneous, depending upon the diligence of the network admin to configure proper values.

- **SNMP Location.** This is the system location of the device acquired from the SNMP agent "sysLocation.0" MIB object value. As above, the value is usually configured by the device administrator, and may be blank or erroneous.

To refresh the screen with the latest information, the user clicks on the navigation tab for the screen, which acquires new data.

If SNMP is not being used by the enterprise, the value of the read community can be set to "none", or the location of the "net-snmp" software removed from the system. This will not affect other functions, and will eliminate the SNMP values from this screen.

Device Associated URLs and Extra Tabs

The administrator can click on the device hyperlink to access the "Device Information" screen, then click the "Edit Device Info" link, and then click "Edit Extra Info and Tabs" on the Device Information editor to specify links and tabs that are displayed on the top-level screen whenever a device IP address is clicked. The operator can specify any URL (and a label) and this URL and label will appear on the Device Information Screen in the upper right.

Additionally, the administrator can create an "Extra Tab Group" (by editing the "config/dev-xtab.cnf" file) and then select this tab group to create a series of up to eight different tabs on the Device Information. This is a more advanced feature, but provides a simple method of "hooking" into the CorreLog navigation for devices. The "config/dev-xtab.cnf" file provides a simple format, as shown below.

```
Group_Name.1_Label      My Tab Name 1
Group_Name.1_URL        /x-cgi/my_script1.php?ip=$ipaddr
Group_Name.1_width      900
Group_Name.1_height     600

Group_Name.2_Label      My Tab Name 2
Group_Name.2_URL        /x-cgi/my_script2.php?dev=$devname
Group_Name.2_width      900
Group_Name.2_height     600
```

As depicted above, the "config/dev-xtab.cnf" file provides the name of a device group, and then up to eight different tab definitions, where each tab has a label, a URL, and an IFRAME size. The first column of the file takes the form (groupname).(num_field), where (num) ranges from 1 to 8, and (field) is either "Label", "URL", "width", or "height".

Multiple device tab group names and tabs for each group can be defined. When the operator clicks on a tab definition (once the tab group has been selected on the "Edit Device Info" screen) the specified URL is launched in an "IFRAME" of the specified height and width. This lets a savvy administrator customize the system with PHP or other URLs of specific interest when viewing a device.

Finally, note that the URLs configured (either in the "dev-xtabn.cnf" file or as the "Associated URL" for the device) can contain the "\$ipaddr" variable, which substitutes for the IP address (or device identifier) in the URL. Also, if a "Device Name" is configured (i.e. the DNS name, or some other value) the "\$devname" variable is substituted for that device name. This permits the URL to access another web server, passing the operant device IP address or name to the third-party system.

Dashboard Screens

One of the main CorreLog applications, and one of the easiest to get started with, is the "Dashboard" facility, which is accessed via the second tab at the top of the screen. The dashboard facility allows the user to select a layout field and then add "Gadgets" that depict the current state of messages, devices, alerts, threads, users, and other real-time system parameters. Multiple dashboards can exist on the system, and users can select their default dashboard (first displayed when the user clicks the "Dashboard" tab) in their "User Preferences".

Dashboards are especially well suited as monitor screens for network operation centers, or as a top-level screen for the entire CorreLog system. In addition to displaying real-time summary status, the user can click down into dashboards to view detailed message data and statistics, making it well suited as a CorreLog navigation tool.

Dashboard Configuration And Layout Files

Each dashboard consists of a single specification file residing within the "CorreLog/dash/config" directory. This configuration file fully specifies the content of the dashboard, including layout file, gadget position, and gadget arguments. These files can be viewed and carefully modified outside of the CorreLog system to modify dashboard content, or to create new dashboards.

The dashboard configuration file is modified via the "Edit Layout" or "Edit Gadget" screens, and also modified whenever the user drags and drops a gadget to a new location. If necessary, the user can repair problematic dashboard installations through direct edits of the dashboard configuration file using a normal text editor.

Each dashboard configuration file references a "Layout" file, which resides in the "CorreLog/dash/layout" directory. These are HTML files containing IFRAME

references and macros, used when a dashboard is displayed. The user can create and modify new layout files to implement dashboards with multiple IFRAME references. Users can also annotate layout files with special HTML markup code (such as to add an organization logo to a dashboard.) This will typically require some understanding of HTML, but even non-programmers should be able to follow the layout file commentary to accomplish activities such as adding multiple new IFRAME references to the layout.

Within the Layout File, the following macros (identified by "@@" characters) are automatically substituted with system information when a dashboard is displayed. The user can incorporate any or all of these macros into the system. Each macro is given an ordinal and increasing integer number, starting with "1", which identifies the windowpane and gadget.

- **@@Pane_N.Name@@** - This macro is substituted for the "Nth" window panel name as configured in the Edit Layout screen of the dashboard. The default value for this macro is a zero length string.
- **@@Pane_N.Link@@** - This macro is substituted for the "Nth" window panel hyperlink for the above name, which permits the user to drill down into a new dashboard. The URL value is configured in the Edit Layout screen of the dashboard, and is useful only if a dashboard name has been specified.
- **@@Pane_N.Editor_Link@@** - This macro is substituted for the link to the dashboard gadget editor for the "Nth" window panel. By default, this is the link that is executed when the user clicks on the "Edit Gadget" image on the dashboard. The user can omit this value, or associated the hyperlink to some other element on the dashboard.
- **@@Pane_N.Exec_Link@@** - This macro is substituted for the link to the actual dashboard gadget. This value is typically referenced in the "IFRAME SRC=" statement. This is the only value actually required in the dashboard, and is needed to display the specified gadget.

There is no actual hard limit to the number of dashboard panes within a layout file. CorreLog comes with three layouts (suitable for 800, 1024, and 1280 resolution monitors) but users can very easily create layout HTML files with many different gadgets. Caution should be taken when creating layout files to prevent the server system or user's browser from becoming too overloaded, which may cause very long rendering times in a browser, however the capacity for most modern systems is quite high, and dashboards with 24 or more panes can be easily created.

Referencing CorreLog Gadgets Externally

Dashboard gadgets can be referenced on external web pages and websites, either requiring a login or not. The simplest way to reference a dashboard gadget is to build a dashboard within CorreLog (to create the configuration file) and then reference the pane number for the dashboard as part of the URL argument that specifies the gadget.

Gadgets are accessed at the URL `/s-cgi/gadget/gadget-name.exe`, and each gadget accepts as an argument a panel number and configuration file reference, in the following form:

```
/s-cgi/gadgets/(gadgetname).exe?(panenum)/(configfile)
```

The above URL runs the gadget program, which gets its detailed setting information from the configuration file and panel number specified.

For example, to add the "Graph-Message-Rate" gadget to an external website (to show an indication of the number of messages received per hour by CorreLog) the user creates and configures the dashboard, giving this dashboard some name such as "External-Dash". The "Graph-Message-Rate" gadget is added (and configured) as the first window panel using the CorreLog editor.

The Graph-Message-Rate gadget is then added via an IFRAME to any website using the following HTML (included on one line.)

```
<iframe  
width=300 height=300  
scroll=off border=0  
src="http://host/s-cgi/gadgets/Graph-Message-Rate-  
Gadget.exe?1/External-Dash"/>
```

Note that the "host" argument above is the hostname (and optional service port number) for the CorreLog server. The value of "1/External-Dash" references the first panel of the External-Dash.cnf file on the system. The result is that the gadget is displayed on the specified web page, using the configuration specified in the "External-Dash" configuration file, panel 1.

Any other gadget in the "CorreLog/s-cgi/gadgets" directory can substitute the executable name "Graph-Message-Rate-Gadget.exe" above, permitting a wide variety of gadgets to be incorporated into the content of an external website.

Note that gadgets can require fairly specific and elaborate arguments, such as message qualifiers, address groups, and other details to make the gadget quite

specific. The gadget will read the "External-Dash" configuration file, will seek the data specified in the named window pane, and display data using the configuration rules specified therein. If the "External-Dash" configuration file is deleted or modified, the gadget on the external web page will reflect that change.

Finally, note that the s-cgi directory and all subdirectories are protected by HTTP authentication, requiring the user to login to the CorreLog system when executing any page containing the above IFRAME before the gadget is displayed. The administrator can simply create a different and unprotected CGI directory (such as "x-cgi/gadget") and copy in the ".exe" file to that location, modifying the HTTP server as needed. This provides a way to bypass the login authentication normally needed to view any CorreLog data, with the specific consent and agreement of the CorreLog administrator.

Graphs Screen

The CorreLog Server incorporates a general-purpose graphing tool, which allows the user to see the hourly or daily message counts. This screen is accessed by clicking the "Reports" and then the "Graphs" tabs at the top of the screen.

The "Search" facility, described previously, provides excellent service and finding data on the system. However, the Search facility (while presenting times) may not be convenient for assessing the time rates, especially over many days. This is the role of the "Graphs" screen, described here.

By default, on entry the "Graphs" screen shows the daily counts of messages for the last 30 days. The counts of all messages are shown. The operator can filter this display by matching a specific device, facility, and or severity. The user can also view data earlier than 30 days ago by adjusting the "Start Date" setting.

Additionally, the "Graphs" screen can show the hourly counts of messages for the last 24 hours. (The user modifies the "Graph Type" setting from "Daily" to "Hourly" and then clicks the "Apply" button.) The operator can view the hourly rates for any daily data on the system by adjusting the "Start Date" setting.

At the bottom of the screen are statistics that summarize the minimum, maximum, average, and standard deviation of all data displayed.

Graph Feature Summary

The graph screen contains the following features:

- **Daily and hourly graphing.** The user can select display a graph of the daily counts, spanning 30 days, or an hourly readout, spanning 24 hours. The user selects the "Graph Type" from the top control bar, and then clicks "Apply".

- **Graph start date.** The user can view the graph data starting at any date on the system. On entry, the graph start date is today's date. If the operator sets the date back, the data starting at the particular selected date is displayed. On "daily" graphs, this means that the user can view the daily counts starting at arbitrary day. On "Hourly" graphs, the user can see the hourly rates for any day.
- **Auto-ranging.** On entry, the graph automatically auto-ranges the top scale of the screen to be the best fit for the data, starting at zero as the minimum value. This shows the best initial depiction of the data, on entry to the screen.
- **Graph scale slider.** The left side of the graph includes a slider switch (discussed in more detail below), which allows the user to easily adjust the maximum range of the graphed data. This is especially valuable if there are various data spikes that may reduce the initial precision of the graph depiction, displayed on entry to the screen.
- **Statistics readout.** At the bottom of the graph is a statistical read out showing the minimum, maximum, average values, and the standard deviation of the data. This reflects all the data depicted on the graph (either the hourly or the daily values, depending upon the setting in the control bar.) It also shows the values after data filters are applied to the screen, as discussed above.
- **Bar hover readout.** If the user hovers over a particular bar, the X / Y values for that bar are displayed. This allows the user to easily inspect the specific data associated with a day or an hour, and then drill down into a particular range of messages.

Graph Scale Slider Bar

It is quite common that on certain days, or during certain hours of the day, there exists large numbers of messages, while for other days or hours there are few or no messages. This means that, initially, the graph often depicts several steep spikes, which causes the graphs to be auto ranged in such a way that all the data is not easily inspected.

The Graph tool includes a slider bar, at the left of the graph, which allows an operator to easily adjust the range to view small data counts. The slider bar is normally at its bottom position, showing the graph range at full auto-range scale. The user can click and drag the bar upwards to reduce the scale, lowering the ceiling, and causing the bars to increase in height. At its maximum position, the slider bar sets the scale to 1 count (which usually results in all bars being clipped, unless there were zero counts for a particular bar.)

As the slider is dragged upwards, the scale axis legend shows the setting in real time. (The y-axis legend values change, showing the new ranges.) When the operator releases the slider, the screen is repainted at its new values. The operator can quickly find a suitable slider position so that the data range of interest is displayed. Note that the slider operates logarithmically, which permits the user to make precise adjustments at lower ranges.

Data Catalogs

The “Search” screen, although extremely useful, does not necessarily provide the overview of the messages, which may be necessary for the operator to adequately assess what devices are busiest, or what message types are received the most. To permit this overview, in addition to providing a high-speed search engine, the Syslog messages are continuously catalog into different groups: Devices, Users, Facilities, and Severities, and Correlation Threads. The purpose of these catalog screens is to furnish a real-time synopsis of message activity, and assist the user in viewing messages related by some common factor.

- **Devices Catalog.** This is a list of all the devices on the system, which can be ordered by time of last message (the default), or by counters (described below.) This screen provides a good overview of the device activity, which is one of the most pertinent views of the data for most managers. The list of devices can become quite long, but the screen can easily be sorted or filtered to view a particular class of device.
- **Users Catalog.** This is a list of all the users on the system, which can be ordered by the time of last message (the default) or by counters or by username. This screen provides a good overview of the users on the network, and all the messages that have been received by CorreLog which reference the user. For example, this screen is highly useful for auditing the logins (and other actions) of a user across multiple managed platforms, essential for security auditing.
- **Facilities Catalog.** This is a list of all the facilities on the system. Like the Devices Catalog screen, the list can be ordered by time of last message (the default), or by counters. The facilities indicate what class or type of message is being received. In particular, CorreLog permits users to define new facilities, discussed later in this section.
- **Severities Catalog.** This is a list of severities on the system. Like the Devices Catalog screen, the list can be ordered by time of last message (the default) or by counters. Severities indicate how severe a particular message is deemed to be. Therefore, this screen indicates the overall health of the managed devices that are communicating with CorreLog.

- **Threads Catalog.** This is a list of special match patterns, configured via the "Correlation > Threads" screen, that arbitrarily correlate messages by simple or complex match patterns containing keywords, logical operators, or other correlation functions. (This type of catalog is detailed in Section 5 of this manual.)
- **Aux Message Catalog.** This screen is slightly different from the above screens. The Aux Message Catalog contains a list of all the messages that have been filtered during the course of a day. This is useful mainly for debug and system maintenance, as a method of determining what messages may not be cataloged because the Syslog process is eliminating them before the Cataloging process sees them.

Each record in the Devices, Facilities, and Severities screens show the IP address, Facility, and Severity of the last message received on the system associated with the catalog item. For example, in addition to showing the most recent device to receive a message, and when the message was received, the "Device Catalog" screen shows the severity and facility of the last message received by a particular device.

Catalog Message And History Counters

In addition to the IP address, Facility, and Severity, each catalog record contains two separate (but related) counters. The first counter (labeled "Today") indicates the number of messages in the catalog since midnight. The second counter (labeled "History") indicates the number of messages in the catalog since the first detection of the catalog item, or since the catalog was last deleted.

The "History" counter can be thought of similar to an odometer reading, in that it is the long-term count of messages for a particular item. This provides an indication of whether a device, facility, or severity has been active over a long time (although it may not have been recently active.) The "Catalog Message List" screen, described below, can clear the "history" counter. Otherwise, the counter continues to increment across platform reboots.

The "facilities" and "severities" counters work with the Correlation functions of the system. Specifically, the user can set alerts on these counter rates, which sends user defined messages back into the Syslog receiver. (This is discussed in detail within Section 5 of this manual.)

The Catalog Message Viewer Screen

The various catalog viewer screens provide summary information on the various cataloged items. The system allow the operator see the list of messages associated with a catalog item by clicking down on the catalog item hyperlink. This displays the "Catalog Message Viewer" screen.

This supports the notion that an administrator may wish to see the detailed information associated with a particularly active severity, facility, user, or device. This is an easy way of getting the list of messages associated with any severity, facility, user or device on the system. The operator can also use the "Search" screen to do this, via a search term, but it may be easier to see these messages in the context of the catalog viewer screen, which also provides overview information not available elsewhere in the system.

The Catalog Message Viewer supports search capability, but all searches are "non-indexed", hence are slightly slower than using the regular "Search" facility. The operator can page through the data via hyperlinks to see the entire results set associated with a catalog item. Messages are always shown in reverse chronological order, with most recent messages displayed first.

The Catalog "Analyze" Function

At the top of each "catalog" screen is an "Analyze" link, which permits the user to drill down into the data to analyze the catalog content in terms of "Devices", "Users", "Facilities", and "Severities". Clicking the "Analyze" function displays a screen that shows catalog items and their frequency of occurrence. Each item is hyperlinked to the set of messages comprising the item. This provides a simple way of viewing catalog data from various perspectives.

Using the "Analyze" function, an operator can quickly examine catalog content to see the list of devices, users, and other items that constitute the catalog. This is particularly useful for viewing "Thread" data, which may consist of diverse messages originating from a variety of devices and citing a variety of users, facilities and severities.

When executing the "Analyze" function, the operator selects the time range over which to analyze data. This "Span" time frame is by default one hour, but the operator can select a larger time frame (such as 10 days.) The maximum number of messages parsed, irrespective of the selected "Span" type, is governed by the "Max Non-Indexed Search" parameter of the "Messages > Config > Parms" screen. Therefore, for a large catalog of messages, the "Span" type selection may not span the actual number of records present in the catalog, especially if the catalog is quite large.

The "Analyze" function is available via the "Analyze" link, and is also available by clicking the "(+)" icon on the top-level catalog screen.

Deleting A Catalog Item

The Catalog Message Viewer contains an "Edit" button, which allows an operator to delete a particular catalog item from the system. (Currently only one edit function is defined, which is to delete the catalog item, with confirmation.)

Note that deleting a catalog item does not actually delete any messages (which are still available for searching via the "Search" screen.) The delete function removes the item from the CO-Catlog.exe program, along with any history acquired by that program. If a message is received that contains the catalog item that was deleted, the entry is re-added.

The user can delete device, facility, or severity catalog items. In practice, it is mainly useful to delete "Device" catalog items (since it is common for a device to be removed from the network, and this "delete" capability is useful for removing the information for that device from the CorreLog Server.) There are several reasons why an operator may wish to delete a particular catalog item.

1. Deleting the catalog frees up disk space. This may be significant if there are a few devices that are generating vast amounts of message information. This may be useful if the device "history" counter is very large, but the message counter is small, and has not incremented in a long time. (This can be viewed by sorting the display via the "History" sort order.)
2. Deleting the catalog sets the "history" counter back to zero. This provides a way to "baseline" a particular device, facility, or severity. For example, the software on a managed device may have changed significantly, so the catalog history no longer reflects the actual use of a managed device.
3. It may be that a particular catalog item is no longer pertinent. Perhaps a device has been deleted from the network, or a user defined facility code is no longer applicable. In particular, this will be needed if the user is experimenting with special facility codes, defining new codes, as discussed in a later part of this section.

Once a catalog item is deleted, the system will re-add the catalog item the next time an appropriate message is received. For example, if the user deletes the 10.1.1.5 device from the "Devices Catalog", all catalog data for the device will be deleted from the system. However, next time that a message is received from the 10.1.1.5 device, the device will be re-added to the "Devices" catalog, and the "history" counter for the device will be set to one (for the one message received since the catalog item was deleted.)

Regenerating A Catalog

When a catalog is first created, it begins recording messages from that point on. Specifically, pointers to the message content (in the main "logs" directory) are derived and tracked. The catalog itself does not duplicate the message, but merely keeps the offset of the message in the main log file.

It may be the case that these pointers become unsynchronized for some reason, such as a line in a log file has been deleted, which affects the offset from the start of the log file. In that rare case, the catalog will become corrupt and unusable.

It may also be the case that the user wishes to build a catalog from the historical data on the system. Because the messages already exist for a catalog, it may be desirable to scan all the messages and rebuild the catalog based upon message history (as opposed to waiting for the catalog to be built in real-time.)

Users can regenerate a catalog via the "Regenerate Catalog" hyperlink, found at the bottom of all catalog viewer screens. This function allows the user to launch (via the web interface) a background process that will scan all the log messages (to some user defined limit) and collect the offsets to each message associated with a catalog. This process may take several minutes or more to accomplish, but will result in the catalog being completely rebuilt.

Because this is a slightly CPU intensive operation, only one catalog can be rebuilt at a time. The user must wait for the catalog to be rebuilt before regenerating a new catalog. The user can also terminate the regenerate process and start another process, in which case the first catalog is not affected.

Regenerating a catalog is especially useful for "Thread" and "User" catalogs, which initially have no associated data when they are created. The "Regenerate Catalog" function permits the user to easily populate these catalogs with historical data.

Note that only users with "admin" system access can regenerate catalogs. This is to prevent data tampering by ordinary users of the system.

Message Filters, Overrides, and Configuration Screens

Previously, this section has discussed the “user level” interface screens, which will be the target of most usage. The CorreLog Server also provides a series of administrative and configuration screens, in the “Messages > Config” tab, used to setup and modify the system.

Note that the operator does not have to configure these screens. (The out-of-box configuration is completely usable, with no modifications required.) However, a user can make adjustments for specific applications using the various “Advanced” screens. In particular, there are two main types of configuration:

- **Message Filters.** The “Messages > Config > Filters” screen allows an operator to specify filters that are applied to incoming messages prior to any other action.
- **Message Overrides.** The “Messages > Config > Overrides” screen allows an operator to specify match patterns that rewrite certain fields of the message, specifically the device IP address, the facility code, and the severity code.
- **Other Configuration Screens.** The “Message > Config” tab also contains other ancillary configuration items, including the ability to configure message forwarding and ancillary settings that affect message processing and reception.

These types of settings require an “admin” type access to setup and modify. Users can inspect the values, and guests are denied all access to these screens. This section discusses the various configuration options available to the user.

Incoming Message Filters

Syslog messages arriving from the network are called “unsolicited”, because the software has no control over whether a message is sent by a device. This contrasts with SNMP “Get” protocol (for example) which is a “request / response” type protocol. This may require the use of filters, which block incoming messages.

There are two types of filters implemented in CorreLog:

- **Filter Match Patterns.** The user clicks down into the “Messages > Config > Filters” screen to view, add, modify, or delete the list of filters on the system. This is a standard CorreLog dialog. The user adds a new filter by clicking the “AddNew” button, and modifies (or deletes) an existing filter by clicking the “Edit” button.

- **De-duplication Filter.** The user clicks down into the “Messages > Config > Parms” screen and modifies the “De-duplicate Message Seconds” value. This parameter specifies the time period during which duplicate messages, received from a network device, are rejected.

The two filters both remove messages before any overrides or catalog of the messages occur. The filtering takes place in the “CO-Syslog.exe” program, and not in the “CO-Catlog.exe” or CO-Systrap.exe programs. The two types of filters are quite different, but both achieve the effect of reducing the load on the system and amount of data that is stored.

Aux Message Files

To avoid confusion as to whether a message has accidentally been filtered, all messages that match a filter pattern are moved to the “Messages > Aux” screen, where they persist for a period of time. The specific files containing these messages reside in the “logs/filt” folder, which contains a file for each class of data that is filtered.

Filtered messages are normally discarded each day at midnight, but can also be archived with the rest of the data as discussed in a later section, as controlled by the “Messages > Config > Parms” screen.

Note that the “Filter” and “Aux” file combination can be an important tool for increasing both the capacity and performance of the server without losing data. Each Auxiliary file reduces the correlation load on the system (because “Aux” files bypass the correlation and alerting functions of the program.) The operator can direct busy devices (such as firewalls) to an auxiliary file where it can be reported upon and archived. This provides a mechanism for “de-cluttering” the message correlation functions and search engine.

Message Overrides

Message overrides operate in a fashion similar to filters above, except that the message data is altered rather than removed from the system. The operator clicks the “Config > Overrides” screen to access the overrides, which are one of the following:

- **Address Overrides.** The operator can set an address override that replaces the address of the incoming message with some other address, especially useful for systems relying on NAT (Network Address Translation.)
- **Facility Overrides.** The operator can set a facility override that replaces the facility code of the incoming message with some other facility, including a “User Defined Facility” described below. Since facility codes

offer a main way of quickly correlating data, this provides a large amount of flexibility associated with data organization.

- **Severity Overrides.** The operator can set a severity override that replaces the severity of the incoming message with some other severity. Since severities are sometimes badly assigned by system vendors, this provides a direct method of adjusting message severities up or down as may be appropriate.
- **Text Overrides.** The operator can blank out certain fields of a message using a text override. This allows the operator to blank out certain fields within a message, such as a credit card number or user password. The text is eliminated from the message and is not logged on the system.

Note that each override can match an address, facility, severity, and text pattern. Therefore, the order in which overrides are created are important: First the message address is overridden, followed by the facility, the severity, and then the text. So a severity override may be applied to a message, based upon a match address, which was first assigned by an address override.

In addition to the above basic override functions, the various override facilities have substantial advanced functions. For example, address overrides can occur dynamically, by parsing the message content. Also, text overrides provide certain features to automatically mask out user names (to promote user anonymity and user privacy.) These functions are typically available under the "Advanced" tab of the above screens. In particular, the address overrides permit the user to change the mode by which address values are displayed, to support DHCP and non-static IP addresses. (See section on DHCP, below.)

Defining New Facility Codes

One of the more unique capabilities of the CorreLog Server is the ability to add new "Facility" codes to the system, via the "Messages > Configure > Overrides > Facilities" screen. This greatly expands the ability to setup match patterns for other sections of the program, and helps organize the data, and make the facility codes more meaningful.

One of the problems with facility codes is that many of them (such as UUCP) are deprecated or are seldom used. It is therefore of great value to be able to define new (and arbitrary) facilities, or map messages to new facilities. This is especially true because some vendors (such as Cisco) select facility codes that have little meaning, such as "local3". The user can modify these facilities on the fly, mapping messages to new facilities based upon message content, severity, existing facility, device, or time of day.

Facility codes are used to organize the data by identifying the particular “facility”, such as an application or computer or organization, which is the source of the message.

For example, the user may wish to define a new facility called “Exchange” that has all the messages related to the Microsoft Exchange Server. Or, a facility called “Http” can be defined and related to an HTTP monitor. Messages that incorporate these facilities can then be more easily cataloged and searched. Additionally, these facilities can be configured in the “Actions” portion of the program to cause specific programs to be run when certain messages are logged.

Creating User Defined Facility Codes

The “User Defined Facility” list is one of the deepest nested screens in the CorreLog Server. To launch this screen, the user clicks on “Messages > Config > Overrides > Facilities” and from the resulting “Facility Override” screen, and then clicks on the “User Defined Facilities List” hyperlink. Although this screen is deeply nested, it is conveniently placed in the “Facility Override” screen, which is the main context for defining new user facilities.

Note that defining a new facility is only useful if the operator overrides some existing facility to be that new facility, hence user-defined overrides ALWAYS need to be configured in the Facility Override Editor screen before they can be cataloged or searched.

Deleting User Defined Facility Codes

When the user defines a new facility code, it appears in the drop down list of the “Advanced Search”, and Correlation screens, as well as other utility screens.

If an operator deletes a facility code, the “Correlation Threads” or “Actions” screens may depend on that code to either catalog data or take actions. In that case, the user-defined facility remains in place, but the operator will not be able to add new instances of the deleted facility anywhere in the system. Deleting the user-defined facility simply removes it from all dropdown menus of the system, but not from any configuration data. The operator can remove the facility instances (if any) from the “Correlation Threads”, “Actions”, or other screens on the system, using the standard “Edit” and “Delete” functions of the screens, as discussed earlier.

Auto-Overriding Device Names / DHCP Support

In addition to overriding device addresses with alternative addresses via the "Message > Config > Overrides > Addresses" tab, CorreLog contains a special "auto-override" function that allows the user to automatically substitute IP addresses with the "Location:" keyword present (by default) in CorreLog Agent programs. This special facility is available via the "Advanced" button on the "Address Override" screen above.

The following two settings are of particular interest.

- **Require Standard IP Addresses.** This setting, on the "Messages > Config > Parms" screen, disables error checking of the IP address values, and allows hostnames or other values to be substituted using the address override function, including the switch below.
- **Auto-Override Agent Addresses.** This setting, on the "Messages > Config > Parms" screen, and also on the "Messages > Config > Overrides > Address > Advanced" screen, causes the system to automatically override the IP address value with the "Location:" value configured for CorreLog agents, so that device IP addresses are substituted for hostnames on all screens.

The "Auto-Override Agent Address" setting is applicable only if "Require Standard IP Addresses = No", and is otherwise ignored. This setting specifically works with CorreLog Agent programs (which have a "Location:" prefix configured.) The setting will automatically override the IP address value with the value configured at the agent for the "Location: " prefix. Since the "Location:" prefix is pre-configured with the device host name (for most agents), this causes the hostname to automatically be substituted by the hostname for the agent program.

Note that setting "Require Standard IP Addresses = No", and "Auto-Override Agent Addresses = Yes", will result in all messages sent by CorreLog agents to be identified by the hostname (or more specifically the value of the %COMPUTERNAME%" setting) on CorreLog screens. This may be particularly useful in a DHCP environment with short expiration leases.

Auto-Overriding Device Names With Parsed Values

The "Messages > Config > Overrides > Advanced" screen provides a special function to override the device name based upon a value parsed from the message content. This is useful if there is a load balancer, single Syslog server source, or other situation where the CorreLog Server cannot necessarily detect the IP address of the originating device based upon the network connection.

For example, if a single Syslog-NG server is the source of all messages, the managed device will normally be the address of the Syslog-NG program, rather than the device that sent the original message to the Syslog-NG collector.

The user can parse the message content to determine the device name using special "Parse" rules, accessed via the "Auto-Override With Parse Specification" button on the "Messages > Config > Overrides > Advanced" screen. The "Parse" rules permit the user to parse the device name from the message content, and subsequently format the device name. The parsed name then appears as the device name within the CorreLog server.

This function, while obscure, may be essential to the operation of CorreLog at those sites where a load balancer, central Syslog collector, or Network Address Translating (NAT) router is part of the management strategy. For detailed notes on this screen, consult the "CorreLog Screen Reference" manual on the "Home" screen of the CorreLog Server, or contact CorreLog support.

Auto-Masking User Names For Data Privacy

The "Message > Config > Overrides > Text" tab permits the operator to override specific text in messages, useful for masking items that may be of a highly sensitive nature (such as credit card info, user passwords, etc.) This facility eliminates certain text from incoming messages upon their arrival.

As a special extension of this function, CorreLog contains a special "auto-mask" function that hides user names from the administrator. (In this case, the user name is still saved on the disk, but is otherwise hidden from operators.) This special function is useful for promoting data privacy, and removing potential operator bias when watching the CorreLog screens for threats.

The "Auto-Mask User Names" function is controlled by various parameters, accessed via the "Advanced" button on the "Messages > Config > Overrides > Text" screen:

- **Auto-Mask User Name Enable.** This setting is set to "Yes" to enable the user masking features. Subsequently, a unique identifier (such as "USER012817") replaces any user name appearing on any screen.
- **User Mask Prefix.** This setting allows the operator to specify the prefix to the user mask. The prefix (and unique user identification number) appears in place of each username on the system throughout the CorreLog Server.
- **User Mask Seed Value.** This setting allows the operator to specify a new seed value / offset that changes all the user identification numbers. This setting is useful if a masked user name is discovered. Adjusting this setting will shift all the user name identifiers to some new value.

- **Audit Account Name.** This setting is the name of the single account that is given permission to see the user names on the system. The operator can specify and CorreLog user name. If there is no user with the specified name, this setting is ignored.
- **User Mask Exclusion List.** This button accesses a list of user names that are excluded from the masking process, such as "Administrator", "root", or other common names that are not necessarily associated with a particular user of the system.

To enable user name masking, the operator sets the value of "Auto-Mask User Names" to be "Yes". Optionally, the operator can also modify the mask prefix (which appears in the place of the masked user name), and establish an audit account that is used to see the unmasked user names. Additionally, the operator can specify an exclusion list of user names that are not masked (such as "Administrator", "root", etc.)

User Name Auto-masking, configured by this screen, does not actually override the data within the message database, but simply masks the user names depicted on CorreLog screens and contained in CorreLog reports. (This is different from the main "Text Override" facility, described earlier, which actually modifies message content before it is written to the disk.) In addition to masking user names on the system, user names are also masked in reports and any notifications and tickets generated by the system.

Message Forwarding

An important function of the server is to forward messages to other locations. This permits CorreLog to operate as a collector in a larger management strategy. The system includes support for four general-purpose message forwarders, and an additional forwarder for each "Aux" file on the system.

To forward messages to another syslog server, the operator simply accesses the "Messages > Config > Forwarding" screen and clicks the "Edit" button. The operator can then set the destination addresses for the syslog messages, and can enable the forwarding using one of several different modes as follows:

- **Enable Relay.** This forwarding mode is mainly useful for forwarding messages to another CorreLog Server. The setting causes the messages to be forwarded to the specified "Send To" destination, where the original IP address of the device is preserved as part of the message. This setting is usually useful ONLY if the destination address is another CorreLog server, since it relies on the special CorreLog message prefix notation to preserve the original device IP address.

- **Enable Relay-ENC.** This forwarding mode is identical to the "Relay" option above, except messages are encrypted using CorreLog basic pseudo-one-time pad encryption, enhancing security. This setting is useful only if the destination address is another CorreLog Server.
- **Enable Forward.** This forwarding mode causes the messages to be forwarded to an arbitrary "Send To" destination, where the original IP address of the device that sent the message is included as the first word of the message. This setting is mainly useful when messages are being forwarded to a third-party SIEM system or data collector.
- **Enable Proxy.** This forwarding mode is similar to the "Enable Forward" setting above, except the message is sent without any modification to the header. No hostname or time value is inserted as part of the message, and no simple formatting (such as removing tab characters) is applied to the message. This setting has application in certain situations such as element managers, or in certain test situations.

When a forwarder is configured for Msg-01 to Msg-04, any message that appears in the "Messages > Search" screen is automatically forwarded. All four of these forwarders are identical, and available to forward messages to four different Syslog Receivers.

When a forwarder is configured for Aux-01 to Aux-16, only the messages that are logged in an Auxiliary file (which has been assigned via the "Messages > Config > Filters" screen) will be forwarded. For example, if the operator wishes to direct all the messages of a firewall to another Syslog receiver or CorreLog Server system, the operator filters these messages to the "Aux-1" file, and then forwards the messages onward by configuring the Aux-1 forwarder.

Additional Forwarding Techniques

The "Messages > Config > Forwarding" screen provides a method of forwarding large amounts of data to another server or syslog collector. Other techniques are available to send data more selectively, and for specific application:

- **Message Forwarding Using Correlation Actions.** The operator can configure one or more "Correlation > Action" programs to selectively forward raw message information to another syslog server using either Syslog or SNMP. These forwarders provide more flexibility than the "Messages > Config > Forwarding" screen, including the ability to reformat messages and match very precise conditions. However, this technique should not be used to forward large amounts of data, since over-use of this function can overload the action queue. Generally, this technique can be used to forward 10 to 50 messages a second maximum, and is best implemented for occasionally occurring events.

- **Ticket Forwarding Using Actions.** The operator can configure one or more "Ticket > Action" programs to forward CorreLog tickets using a variety of techniques, including Syslog, SNMP, and TCP Tunnel programs. (These are discussed in more detail in Section 6 of this manual.) The same flexibility and limitations exist for ticket forwarding as message forwarding, described above. This type of forwarding is especially useful in two-tier management strategies, as discussed below.

Forwarding to Support Multi-Tier Operation

A main reason for forwarding messages is to support multi-tier management, as a method of scaling or organizing data. Any site that is interested in multi-tier operation must necessarily consider the forwarding capabilities and strategy of the program.

One common multi-tier strategy is to send all syslog messages to a central collector, which then keeps some of the data locally, and send other parts of the data to a higher-level collector. In the case of CorreLog Server, the operator simply configures the correlation rules and alerting for a single CorreLog site, and then forwards ticket information to a higher-level manager, which may or may not be another copy of CorreLog Server. This permits massive scaling of information, where multiple servers exist in an enterprise.

In this scenario, each copy of CorreLog is responsible for a particular class of device, department, user set, or other logical partition of data. Correlated results are forwarded up to a top-level manager.

This strategy supports massively scaling of CPU and disk resources. Additionally, this type of distributed system keeps the data segregated (which may be an important security concern) and reduces network congestion and traffic by limiting the distance over which messages must be transmitted.

Note that CorreLog is especially well suited for this management strategy since it is quick to install, completely web based, and can co-exist with other software systems (reducing the need to buy hardware or install appliances.) This type of strategy can easily accommodate applications requiring many 100 thousand events per second, potentially millions of devices and users.

Other Interactive And Administrative Screens

The CorreLog Server provides several other interactive screens, mainly useful for high-level system configuration, or for use under the direction of customer service. These screens include the following:

- **Color Editor Tool.** This tool is accessed via the “Messages > Config” screen, and allows the user to set the colors associated with severities displayed in various locations within the system. The user can select colors for both the text and the background (with the aim of permitting easy contrast between these.) This screen requires an “admin” type login to modify.
- **Parameter Editor Tool.** This tool is accessed via the “Messages > Config” screen, and allows the user to set various parameters of the system that affect the performance of the Syslog processes and web interface. This screen requires an “admin” type login to modify.
- **Sendlog Tool.** This tool is accessed by clicking the bottom link of the “Messages > Config > Parms” screen, and allows the user to send a Syslog message to the CorreLog server, useful for test and verification of system settings, and to verify that the system is running correctly. This screen requires an “admin” type login to execute the tool.
- **System Info Tool.** This tool is accessed by clicking the bottom link of the “Messages > Config > Parms” screen, and allows the user to view certain system constants. This screen is mainly useful for sending data to customer support. This screen requires an “admin” type login to view.

The above screens, while useful, are generally accessed rarely. The exception may be the “Sendlog” tool, which furnishes the ability to send a Syslog message to the CorreLog Server. This tests that the Syslog server process is running, but can also be used to test filters, overrides, and actions.

Data Archiving Functions

The CorreLog system performs data archiving by compressing log files each night, and generating MD5 message digests. These files are stored in the directory specified on the "Messages > Config > Parms" screen, by default the directory "archive" within the CorreLog root directory. The "CO-maint.exe" program is responsible for this process, and creates an archive file each night as follows:

At midnight, the system checks each file in the "logs" directory, and for each file that does not have a ".gz" file in the "archive" directory, an archive file is created in Gzip format. Therefore, if the system is not running at midnight on a particular

night, the archive for a file will be created on the subsequent night. This permits the system to be shutdown during midnight with no loss of data, provided that the system is up at least once at midnight while an active log file is still on the system.

Restoring Archive Files

To restore a data archive, the user can unzip the file back to the folder, creating a new entry in the "logs" directory of the system. Then the user can index file so that it appears in the "Search" screen. This can be accomplished via the "RestoreArc.exe" program, residing in the "system" directory.

For example, to restore an archive file, created on January 31st, 2014, the user executes the following at a command prompt, from within the CorreLog "system" directory:

```
RestoreArc.exe 2014-01-31 -overwrite
```

The above command accepts one or two arguments: The first argument is the date of the archive in YYYY-MM-DD format, and the second optional argument will overwrite any existing log file (if it exists.) The user can execute the above command at a command prompt, or automate restoration of archive files via a batch execution file.

Once the archive is restored, you can search for data via the "Search", "Graph" and "Query" function. HOWEVER, restoring the archive file DOES NOT necessarily rebuild the index files for the catalogs (i.e. the "devices", "users", "facilities", "severities", and "threads" screens.) To explicitly rebuild these files (which may take some time) you should explicitly rebuild the index files with the following command:

```
RestoreArc.exe 2014-01-31 -index
```

The above command will make the data accessible to the various catalog files on the system. Depending on the amount of data to index, this may take only a few minutes, or may take several hours.

Note that when an archive is restored, it will be deleted at midnight (if the date of the archive file is greater than the "Keep Online Data" setting of the "Messages > Config > Parm" screen.)

Other Ways Of Using Archive Files

Archive files can also be searched using the Reports > Query function. This screen (which is normally used to search online log data) can also search archives, unzipping these files automatically for search. The user can also search

archives by establishing an "External" directory (using the "Messages > Config > Parms" screen) and then using the "External" function in the Reports > Query screen.

Finally, in special cases, archives can be restored via the web interface by regenerating a catalog via the "Regenerate Catalog Information" link found at the bottom of the "Correlation > Threads > Edit" screens. Specifically, the user can click on the "Regenerate Catalog Information" link, then click the "Edit" button (on the regenerate screen) then select the "Regenerate and Restore Catalog From Archive" button. This will rebuild a catalog from an archive, and optionally restore the archive to the system.

Data Archiving Message Digest Functions

In addition to archiving the file, the system creates a "digest" for the file, which includes an MD5 checksum for the archive file, and other information. The digest file resides in the "archive/digests" folder, with the same name as the archive except with a ".txt" suffix.

The MD5 checksum, contained in the "digests" folder, can be used to verify that the file has not been tampered with. To enhance the forensic capabilities of the system, the MD5 checksum is further encrypted to create a "Security Code". This security code provides absolute tracking of information, and can only be deciphered by the vendor. (This might be useful in conflict resolution, for example, or in detailed forensic analysis of the log file information.)

CHKSC Program – Checking Archive Security Codes

The "chksc.exe" program is included with CorreLog, in the system directory, to allow the user to determine whether any archive has been tampered with. This program decrypts the security code for a particular digest file, and displays an indication of whether the archive has been tampered with.

For example, to detect whether an archive, created on January 31st, 2014, has been tampered with, the user executes the following at a command prompt, from within the CorreLog "system" directory:

```
chksc.exe 2014-01-31
```

The above command accepts either an archive date in YYYY-MM-DD format, or can also accept the pathname to a message digest function, including the ".txt" extension for the file.

Message digests are found in the CorreLog "archive" directory, and contain information about each archive. Because the message digest is encrypted with a one-way algorithm, the user cannot modify the message digest without detection,

creating a verifiable "chain of evidence" regarding each archive, proving that the archive file has not been tampered with since it was created.

Archiving Filtered Data

Normally, filtered messages are retained for a single day and then discarded at midnight. However, may be useful to retain and archive all filtered messages. This may be useful for audit purposes or for data completeness.

If the user has sufficient disk space to archive all filtered messages, archiving of filtered data can be enabled via the CorreLog "Messages > Config > Parms" screen. Setting the "Keep Archived Data" value to "Yes" will cause the CO-maint.exe program to create a new "filt" directory in the archive directory, which will contain timestamped files of filtered messages. These files will have a direct association with the main archived files, and will be kept in the archive for the "Keep Archived Data" duration.

Note that no utilities are provided for restoring these archive files (as exists for normal archive files via the "RestoreArc.exe" program described earlier). To examine the filter data archives, the files can be unzipped to a folder using the "gzip.exe -d" program (to create a text file), and then viewed with command line tools or other programs if extraordinary circumstances warrant that type of detailed forensic research.

Change Audit Capability

The CorreLog system provides a built-in audit capability which tracks all changes made to the system. By default, when the user makes a change, this is logged with a severity of "debug". The actual severity (including a "disabled" severity) is adjustable on the "System > Parms" screen.

Each message includes the name of the user, a description of the change made, and the IP address of the user who made the change. The messages can be used to track changes on the system that may affect message reception and correlation.

Correlations, Actions, Tickets, And Reporting Functions

The functions that have not been discussed in this section are the automation functions of the system, controlled by the "Actions" screen in the "Correlation" section of the program, the "Tickets" screens, and the "Reporting" function.

These functions permit users to watch for specific patterns in the system, and take specific actions when these patterns exist. These functions also permit automatic generation of summary report data on a nightly basis, and publication

of this data using RSS protocol. The CorreLog correlation and automation functions are subtle and very distinctive.

These topics are worthy of their own sections, and are discussed in detail within Section 5, through Section 8 of this manual.

Section Summary And Additional Notes

1. The “CO-Syslog.exe”, “CO-Systrap.exe”, “CO-Catlog.exe” and “CO-Svc.exe” programs run as persistent background processes, and should always be displayed in the Windows Task Manager for the system to be operating properly.
2. The “CorreLog/logs” directory contains the Syslog files, written in chronological order, where each day’s data is held in that directory within its own time stamped file
3. The “CorreLog/logs” directory can be relocated to a different location if disk space is a concern. This setting can be made at the “Messages > Config > Parms” screen.
4. The CorreLog search engine uses the GenDex program to create index files. These files are created periodically, maintained, and stored in the “CorreLog/logs/dex” directory.
5. The “Search” facility includes an “Advanced Search” screen, accessed via a link at the top of the display. This permits users to make more sophisticated searches, including partial matches and exclude keywords.
6. The CorreLog Server displays all devices as IP addresses. More information on an address (including SNMP information, if available) can be obtained by clicking on a device IP address hyperlink.
7. The “Graphs” facility permits data to be graphed on a daily or hourly basis, and incorporates various features to filter and select the graph data, auto-range and adjust the display scale, and display statistics.
8. The CorreLog cataloging function provides a way of organizing data into meaningful chunks and categories. The operator can see the time that a catalog was updated, the number of messages cataloged since system startup, and the historical count of all messages.
9. Catalog data is stored in the “CorreLog/catalogs” directory. This directory can be relocated if disk space is a concern. This setting can be made at the “Messages > Config > Parms” screen.

10. On a "Catalog" screen, the operator can click on a catalog item hyperlink to view all the messages associated with that catalog item.
11. The user can filter incoming messages via the "Messages > Config > Filters" screen. Filtered messages for the last day can be seen in the "Messages > Aux" screen.
12. In addition to these filters, the CorreLog Server includes a "de-duplicate message" setting, which prevents identical messages from being logged within a certain time frame. This setting is available via the "Messages > Config > Parms" screen.
13. The user can override the device name, facility, or severity of messages before they are stored and cataloged. This provides a large degree of control over these values (necessary because of inconsistency that is sometimes found in Syslog messages, as defined by the original Syslog message designer.)
14. The user can define new facility names, which provides the unique and powerful capability with regard to search and cataloging of message data. These facility names appear in most (but not all) of the CorreLog screens, permitting the user to filter, catalog, and graph specific user defined facility codes.
15. The user can forward messages to another Syslog collector or CorreLog server using the "Messages > Config > Forwarding" screen, as well as both Correlation > Actions and "Ticket > Actions" screens. This is especially important for implementing a multi-tier management strategy.
16. Various system parameters are available from the "Messages > Config > Parms" screen, including the amount of time to keep data on the system, as well as access to detailed system information and a "Send Syslog Message" tool.
17. Data is archived in the "archives" directory of the installation, in Gzip format. The user can restore these files, and use the message digest files for forensic analysis.
18. Archives can be restored via the "RestoreArc.exe" command line utility, residing in the "system" directory of the CorreLog installation. This program will unzip and re-index archive files by their date identifier.
19. Archive integrity can be checked via the "ChkSC.exe" command line utility, residing in the "system" directory of the CorreLog installation. This program will decode the "Security Code" included as part of the message digest file, and verify that the archive file has not been tampered with.

Section 5: CorreLog Correlation

Previous sections elaborated upon the basic usage of the CorreLog Server to view and inspect message data. This new section changes the focus of the discussion to identify one of the most distinctive and powerful features of the CorreLog program, which is how to "correlate" messages into groups and patterns, and how to take action when these patterns are detected.

Correlation is defined simply as finding arbitrary associations between messages. The term can be used broadly, or quite specifically. In the case of the CorreLog Server, tools are provided to perform "semantic" correlation (as opposed to a statistical correlation) to detect the specific meaning of events, in real time, so that consequential real-time actions (such as sending notifications, or executing other actions) can reliably occur. Additionally, the CorreLog correlation feature is very useful in reducing the amount of incoming data to provide a clear picture of what is happening with managed devices.

The Syslog portions of the program, discussed in previous sections, performs minor correlation functions (such as cataloging messages by device, facility, and severity.) The "Correlation" facility of the CorreLog Server, described in this section, takes this basic correlation to a much higher and more sophisticated level, allowing users to detect complex patterns and trends.

More correlation features are provided in the "Advanced Correlation System User Guide" manual, which is an embedded manual at all CorreLog sites. That manual is available in the "s-doc" directory of the distribution, and contains valuable supplemental information.

Basic Correlation Components

Five basic "atomic" components comprise the CorreLog correlation facility. These components each play a role in detecting a sequence of real-time event messages, and each work independent of each other, or in conjunction with each other, using the Syslog facility as a connection.

- **Threading Component.** The purpose of the "Correlation / Threads" component is to categorize event messages into "Threads" based upon simple or complex expressions. This gives the user an ability to define classes of related messages. Importantly, this facility keeps track of the "counts" associated with each class of event.
- **Alerting Component.** The purpose of the "Alerts / Counters" component is to detect when thread counters (or other counters on the system) increment above or below a certain rate. This detects when a particular class of events is occurring more frequently, indicating a problem or change of system state. The alert component is concerned with message "rates" and "rate correlation". When a threshold is violated, a Syslog message is sent back into the system where it can be further correlated by the threading component or one of the other components.
- **Trigger Component.** The purpose of the "Correlation / Triggers" component is to detect when a particular event occurs, and then open a "gate" that enables further threading and actions. This is a way of applying a "context" to an event stream. Whereas the Thread and Alert components are mainly concerned with counters, the Trigger component is mainly concerned with flags, each with an expiration time.
- **Pattern Detection Component.** The purpose of the "Alerts / Patterns" component is to detect specific patterns of message states. The screen allows the user to define combinations of triggers that, when satisfied, causes an alert to be generated. In some ways, this is similar to the "Alerts" screen, except offers a different perspective. (The "Alerting Component", described above, is concerned more with message rates, whereas the "Pattern Detection" component is concerned with the state of the system based upon particular messages that have been received.)
- **Action Component.** The purpose of the "Correlation / Action" component is to actually cause some action to occur in response to a message. This action is triggered by any message, but is typically triggered by an Alert, and used in conjunction with a Trigger (described above) so that actions occur only when specific messages correlate at specific rates as qualified by previously received messages. For example, an action might occur if a particular class of event occurs above a certain counts per minute (as

detected by the "Alerts" component", and only if some previous event has occurred at a specific time.

The above components work together to furnish powerful correlation capability. Each of the above components is an atomic component that can be used in a larger correlation strategy. Users can configure Triggers, which enable thread counters, which generate alerts, which are further correlated by threads, which trigger actions, which are enabled by other Triggers, etc.

Correlation Rules And Rule Systems

The configuration data associated with a correlation component is commonly referred to as a "correlation rule". (This term is used throughout this section.) A "correlation rule" is one or more configuration items associated with "Threads", "Alerts", "Triggers", and "Actions". Various correlation rules combine to form a correlation system. The atomic components of the correlation system can be wired in a variety of ways, including (but not limited to) the following:

- **Syslog To Thread Connection.** A thread can be used to count the number of received messages that meet specific criteria. As messages are received, they are compared to match expressions associated with each thread. If a match occurs, the thread counter is incremented.
- **Thread To Alert Connection.** A thread can be connected directly to an alert by setting one or more alert thresholds on a thread counter. The alert adds a threshold to a thread counter (or other counter) and when the threshold is violated, a Syslog message is generated.
- **Alert to Thread Connection.** The output of an alert can be further correlated by a thread, so that a counter is incremented when one or more alerts are triggered. The Alert message is simply another Syslog message. Therefore, it can be threaded and counted in a fashion identical to any other Syslog message.
- **Alert to Action Connection.** The output of an alert can be used to trigger the execution of an action program. The action program is triggered by a match expression applied to an incoming message. That message can be a normal Syslog message, or the Syslog message generated by the "Alert" component. Since the user can define the alert messages and action keywords, this provides a high degree of control.
- **Syslog To Trigger Connection.** A Trigger can be used to flag the occurrence of Syslog messages that meets specific criteria. When a message is received, it is compared against a series of match expressions. If the expression is matched, the Trigger value is either set

(with an expiration time) or cleared. This provides a way of establishing a message context.

- **Trigger To Thread Connection.** A Trigger can be used to qualify and gate the update of a Trigger counter. Specifically, in addition to qualifying incoming messages by match expressions, the operator can specify a Trigger name and state. The counter is only incremented if the Trigger is set (or cleared) by a previously received message.
- **Trigger to Trigger Connection.** Combinations of trigger states can be used to generate alerts via the "Patterns" screen. By monitoring the state of various triggers, specific patterns of messages can be easily detected, such as when events A and B and C have recently occurred on the network. When patterns are detected, a Syslog message is sent back to the system, where it can be threaded and counted in a fashion identical to any other Syslog message.
- **Trigger To Action Connection.** A Trigger can be used to qualify and gate the execution of an action program, the same as with a thread. The action is only executed if a specified Trigger is set (or cleared) by a previously received message.

A more detailed matrix of connections is provided below.

Connect From \ To	Threads	Alerts	Triggers	Actions
Threads		Alerts occur when specific thread count thresholds are exceeded.		
Alerts	Alert messages can be further correlated by threads.		Alerts can set and clear Triggers.	Alerts can trigger action programs to be executed.
Triggers	Triggers can enable/disable thread counters	Alerts occur when specific Trigger count thresholds are exceeded.	Combinations of triggers can send alerts.	Triggers can enable/disable execution of action programs.

The above connections permit sophisticated correlation systems that consist of multiple threads, triggers, and alerts that all qualify the execution of an action program on the system. Additional connections (such as the Trigger to Alert connection) are also possible, and may be employed in certain situations.

For example, a moderately complex correlation might consist of sending e-mail when more than 10 error messages occur within five minutes after any user logs into a system. This would be accomplished by creating a "System Startup Login Errors" thread, and "Login" Trigger. A user could then configure an alert when the number of "System Startup Login Errors" exceeded 10 per five minutes, which would trigger an action to send e-mail.

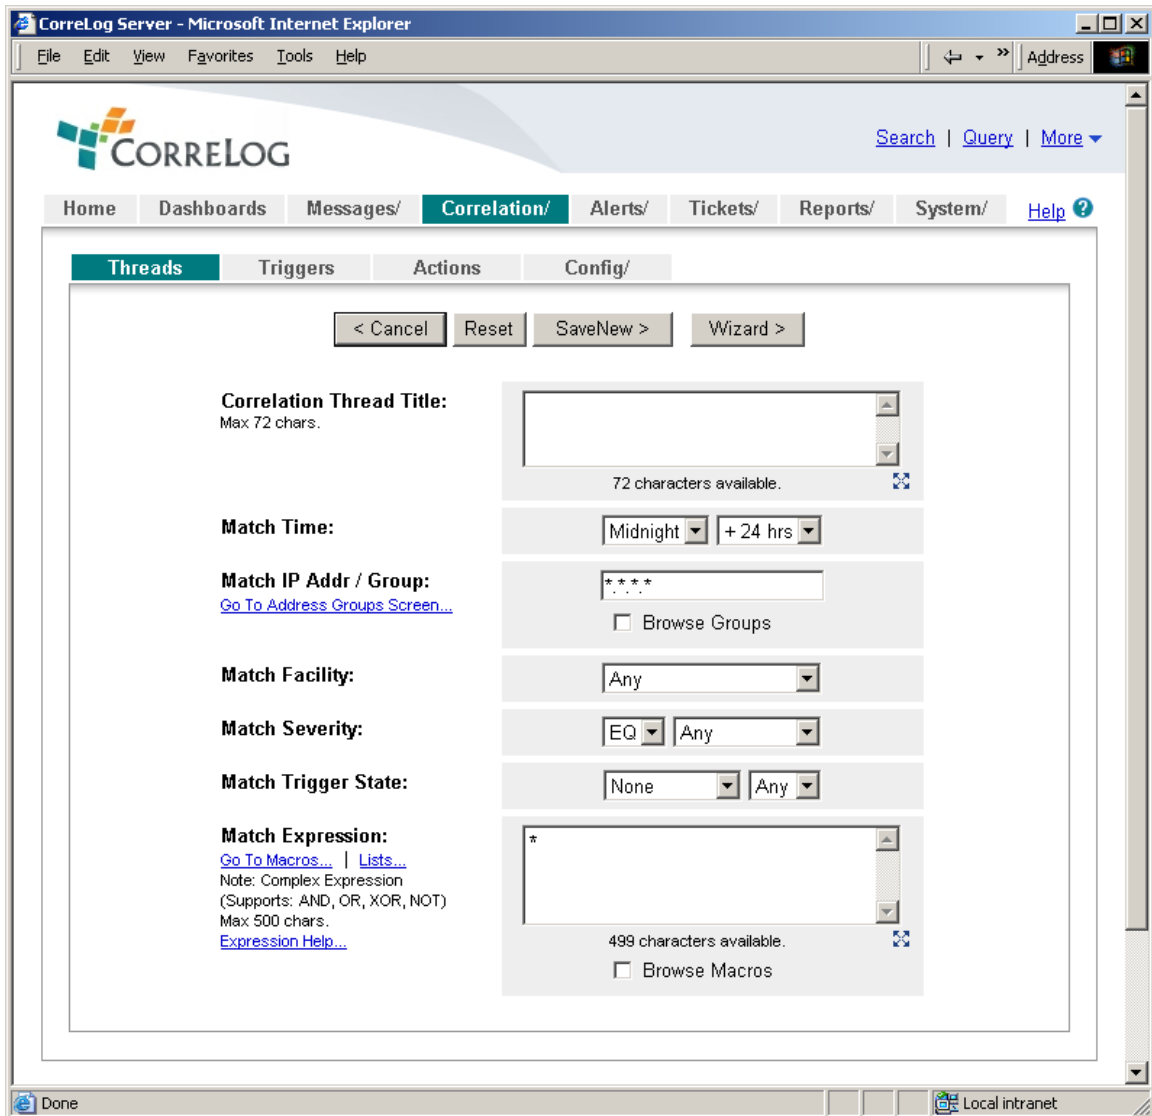
Configuration Of Correlation Rules

Correlation rules are configured in a fashion similar to other screens discussed previously. The "Threads", "Alerts", "Triggers", "Patterns", and "Actions" screens all share common screen controls. Users add a correlation rule by clicking the "AddNew" button. Users can view, modify or delete a correlation rule by clicking the "Edit" button at the left of the correlation rule.

To assist in the activities associated with configuring correlation rules, all of these screens employ "Wizards" that will guide the user through the process of configuring a new rule and configuration item. The configuration "Wizard" is accessible on the top-level screen of each correlation component, and is also available on the "Add New" screen. Each wizard provides narrative, explanation, and help in configuring the data item.

Correlation components work together, and have "Dependants" that are defined elsewhere in the system. The user is blocked from deleting any correlation item that has a dependant, unless the dependant is deleted first. (For example, the user cannot delete a "Thread" if that thread is used in the "Reports".) The user can view any dependants associated with any correlation item by first drilling down into the correlation item (via the "Edit" button) and then clicking the "View Dependants" hyperlink at the lower left of the resulting screen.

To illustrate a typical correlation configuration activity, consider the "Add New Thread" screen, depicted below. This screen is typical of the various Add and Edit screens of the correlation system.



The above screen is similar to the dialogs found in other places within the system. However, several important differences should be noted:

- Configuration Wizard.** On all correlation screens, a "Wizard" button is provided to assist the user in the configuration of the correlation rule appropriate for the particular component. The user may either fill the information into the "Add New" dialog directly, or click "Wizard" for detailed assistance.
- Match IP Address / Group.** The user can specify an IP address, or an address group (defined in the Correlation > Config > Address Groups screen.) Whereas the screens in the "Messages" tab all accept simple IP addresses, the "Correlation" screens accept more complex expressions, including wildcards, and lists of included and excluded addresses.

- **Match Expression.** Unlike other CorreLog screens (such as the Syslog filter and override configuration screens), where a "Match Keyword" must consist of a single word or wildcard, the Correlation screens all permit complex expressions, discussed in a later paragraph. This match expression can be a simple keyword or wildcard, or may be a more complex expression supporting "and", "or", "xor", "not", and containing field matches and parenthetical nesting of expressions.
- **Match Trigger State.** In addition to the various match qualifiers of other screens, such as "Match Time" and "Match Severity", the "Threads" and "Actions" components contain the ability to match a "Trigger" state of either "Set" or "Clear". This furnishes a special ability to "gate on" or "gate off" the correlation counters and actions based upon a previously received event message, which establishes a context.

Except for the above elements, Correlation screens operate in a way that is similar to other screens. The easiest way to get started is to click the "AddNew" button at the top right of these screens, and then click "Wizard", which will guide the user through the process of adding a new correlation rule to the system.

Match Expressions

CorreLog match expressions can be simple keywords, keywords and wildcards, or can be more complex expressions incorporating "and", "or", "xor", and "not" conjunctives. Expressions can contain parenthetical nesting to change the order of evaluation.

- **Case Insensitive Matches.** All matches are case insensitive. For example an expression such as "test" will match "test", "Test", "TEST", or any combination of upper and lower case characters.
- **Partial Matches To Keywords.** Matches can be partial. For example, an expression such as "fail" will match "failure" and "nofail" or any occurrence of the keyword "fail" within the message.
- **Wildcard Matches.** Any keyword can also contain the wildcards "*" to match zero or more occurrences of a string, or "?" to match a single character, or "^" to match the beginning of the line.
- **Literal Matches.** To match a literal string, the string must be enclosed in quotes. For example, "log off" will match any occurrence of "log off" in a string, including the space character. Single or double quotes can be used. Also any single keyword or wildcard expression can be quoted, (although this is generally not necessary unless the keyword or wildcard contains space characters.)

- **AND Operator.** To require various keywords in the message, use the AND operator. For example, the expression "login and user" will match only if the keyword "login" and the keyword "logon" both occur in the message.
- **OR Operator.** To require a selection of keywords in the message, use the OR operator. For example, the expression "login and logon" will match a message that contains the keyword "login" or the keyword "logon".
- **XOR Operator.** To require either of two keyword, but not both, use the XOR operator. For example, the expression "login xor logon" will match a message that contains the keyword "login" or the keyword "logon", but not both.
- **NOT Operator.** To negate a keyword, prefix it with the NOT operator. For example, the expression "not login" matches any message that does not contain the keyword login. Similarly, the expression "login and not system" matches a message that contains the keyword "login", and does not contain the keyword "system".
- **Parenthetical Nesting.** More complex expressions can be constructed using parenthesis to determine the order of evaluation. For example, the expression "(login or logon) and (system or windows)" matches the message "login system", "login windows", "logon system", and "logon windows".
- **Field Matches.** Parenthetical terms can include matches to specific fields, such as "(A in \$5)" where the keyword "A" must appear in the fifth word of the message portion of the input line. CorreLog permits both the "in" and "eq" operators in field matches, where "(\$5 eq XXX)" requires the fifth word of the message to precisely match "XXX", and "(ZZZ in \$12)" requires the keyword ZZZ to appear in the \$12 word of the message.
- **Parsed Field Matches.** Expressions can contain special "parser" functions that match a particular field ONLY if the field has a property such as the field is an e-mail address, number, url, geographical location, etc. Parse functions are documented in the reporting section, but can also be applied to match expressions. (See later section herein.)

Match Expression Examples

Together, these rules permit a wide degree of flexibility in matching input lines. Some examples of these rules are shown below.

Hello	Match any occurrence of "Hello" or "hello" or "HELLO" in the input line.
Not Hello	Match any line that does not contain an occurrence of "hello". As before, all matches are case insensitive.
"Hello World"	Match any occurrence of "Hello World" in the input line, including spaces. The pattern must be enclosed in balanced single or double quote marks, because the pattern contains a space.
Hello*World	Matches "Hello World", and "Hello There World", and any other input line containing "Hello", followed by any number of characters, followed by "World". Since the keyword contains no spaces, quote marks are optional.
Hello or World	Matches any line that contains the keyword "Hello", or the keyword "World", or both, in any order.
Hello xor World	Matches any line that contains the keyword "Hello", or the keyword "World", but not both.
Hello and not World	Matches any line that contains the keyword "Hello", except for those lines that also contain the keyword "World"
not (Hello or World)	Matches any line that does not contain the keyword "Hello" or the keyword "World".
(not Hello) and (not World)	Identical to the above. Match any line that does not contain the keyword "Hello" and also does not contain the keyword "World"
(Hello eq \$7)	Matches any line where the seventh word of the message, delimited by white space is "Hello".
not (World in \$5)	Matches any line where the fifth word of the message, delimited by white space, does not contain the keyword "World". (Note that the "in" operator permits partial matches, but otherwise operates similarly to the "eq" operator.)

Parse Expression Examples

The user can match any particular field via \$N notation as described above. Other more sophisticated parsing functions exist that permit a particular field to be matched only if the field has some property such as being an integer value, or being a URL. CorreLog provides multiple parsing functions that are documented in the Section 8 of this manual.

\$integer(3)	Matches any line where the third word of the line is an integer value, or contains at least one integer number.
\$alpha(4)	Matches any line where the fourth word of the line is an alpha value that is not an integer, and not a punctuation mark. For example, the line is matched if the fourth word of the line is "XYZ", but not "123".
\$basename(5) eq readme	Matches any line where the fifth word of the line has a basename of "readme". For example, the line is matched if the fifth word of the line is "readme.txt" or "readme.hlp", but is not "help.txt".
\$suffix(5) eq txt	Matches any line where the fifth word of the line has a suffix of "txt". For example, the line is matched if the fifth word is "readme.txt" or "help.txt" but not "readme.doc".
correlog in \$url(5)	Matches any line where the fifth word of the line is a URL beginning with "http://" or "https://" or "ftp://", and the parsed url has "correlog" in its content.
\$ipaddr(6)	Matches any line where the sixth word of the line is a standard IPv4 address.
\$geo(6) eq us	Matches any line where the sixth word of the line is a standard IPv4 address whose associated country code is US (i.e. match a US IP address.)
\$email(7)	Matches any line where the seventh word of the line is a standard e-mail address in the form x@y.z.
error in \$delim(pipe, 7)	Match any line where the seventh field, delimited by a pipe character, contains the keyword "error".

Additional parse specifications exist, documented in Section 8 of this manual. See "Special Parsing Rules" in that section for a list of all the different parse specifications and their general usage.

Macro Expressions

Correlation expressions can become quite complex, and are often reused in different locations. To facilitate the configuration and maintenance of these expressions, CorreLog employs a "macro" facility.

Macros are defined in the "Correlation > Config > Macros" screen, and correspond to any expression. The user defines a macro name, the expression, and then uses that macro in the "Match Expression" field of the "Threads", "Actions", or "Triggers" screen.

Macros have a special naming convention: each macro is in the form @@name@@, where "name" is the name of the macro. When the user enters a macro name, CorreLog substitutes it for the value of the macro, which can be a simple or highly complex expression. The user can modify the macro value in the "Config > Macros" screen to effect an immediate change to the system, promoting easy maintenance of correlation rules.

Note that macros cannot be nested, but several macros can be used in a single expression. For example, the user can define a macro named "@@logins@" and a macro named "@@database@", and then enter the correlation match expression of "@@logins@" and not "@@database@" to match any message that is a login and not a database.

Macros can be entered only in "Correlation" screens, and cannot be used as filters or overrides in the "Messages" screen.

List Macro Expressions

"List" type macros are a special form of macro, which can contain large numbers of items that are logically "orred" together. A "List" type macro can be used anywhere a regular macro is used, to create blacklists, whitelists, or other constructs. Unlike "Macros" (which are limited to a small number of characters, such as 500 characters or less on most systems) the "List" type macro can contain 1000's of different items.

Lists are defined in the "Correlation >: Config > Lists" screen, and contain items that are compared to each word in the incoming message. (This function uses a high-speed binary search, hence the correlation engine can quickly process hundreds or thousands of list items.) To match any item in the list, the operator simply specifies the macro @@name@@ as part of the expression. To exclude any item in the list, the operator simply specifies "not @@name@".

List macros can be used in conjunction with regular macros, or with other expressions. For example, the operator can specify a match expression such as

"@@list1@@ and @@list2@@ and not (@@list3@@ or discard)". Note that, using this technique, lists can be joined, but lists cannot otherwise be nested.

Address Groups

Address groups are particular type of macro, similar to the above macros. Address groups have the same naming convention in the form @@group@@, however an address group cannot be used as a macro, and a macro cannot be used as an address group. Specifically, address groups are used only in the "Match IP Address / Group" fields of the "Threads" and "Actions" screen.

Address groups are defined in the "Correlation > Config > Address Groups" screen, and consist of lists of IP addresses to include and exclude in the group. These addresses can be fully qualified, or can contain wildcard characters, such as "10.2.*.*", which matches all devices on the 192.168 Class-B subnet.

Note that this provides an effective way of managing messages by device group. However, it should be noted that address groups can easily be overused, and that managing messages by content (as opposed to device group) is typically more powerful, and a more accepted technique. When defining and employing an address group, the user should consider whether this is really the optimal way of managing messages, given the objectives of the organization. For example, it is probably not necessary to create an address group called "Oracle Machines", since the match expression value will target Oracle messages irrespective of the IP address that sends the message.

Thread Groups

Thread groups provide another organizational level to the system. Unlike "Macros", "Lists", and "Address" groups, the "Thread Groups" are not used in correlation strictly, but rather group correlation threads into similar patterns (such as "Windows" type threads, or "Firewall" type threads.)

Thread Groups are match expressions that are applied to thread titles. The user can configure thread groups via the "Correlation > Config > Thread Groups" screen, and can select thread groups via a drop-down menu at the top of the "Correlation > Threads" screen. The user can also view the status of correlation threads via the "View Threads" link on the "Correlation > Threads" screen, and can select a particular "Initial Thread Group" via the "System > Prefs" (so that the thread group of interest is automatically accessed when the user selects the "Thread" screen.

Some important considerations apply when a thread group is displayed: If the user attempts to add a thread or rename an existing thread, the specified thread MUST match the thread group, or an error indication is generated by the system. In this case, the user should either set the "All" thread group and repeat the

action, or should include a keyword in the thread title that matches the Thread Group. For example, if the "Windows" thread group is selected (which matches the keyword "Windows") and the user attempts to create a thread called "Error Messages", the user will see an error message. The user simply clicks the back button to return to the thread editor, and renames the thread "Windows Error Messages" to allow the new thread to be created.

Alerting Functions

The Alert function, which feeds back data into threads and triggers (and also optionally open tickets on the system) requires configuration of two specific parameters: the alert threshold, and the sample interval. To simplify the configuration of these items, the user can click on the "View Counter Threshold Hints" hyperlink available on the Alert "Edit" screen, and the "Add New Alert" wizard. This link launches the "Message Statistics" screen, depicted in Section 2 of this manual. The "Message Statistics" screen displays "Threshold Hints", which suggest values for Alert settings as follows:

- **Info Alert Threshold Hint.** This value is the average value for the counter over the specified test interval, plus one standard deviation from the average. Statistically, this threshold will be met approximately 35% of the time during the span of measured counts. The minimum value is 1.
- **Notice Alert Threshold Hint.** This value is the average value for the counter over the specified test interval, plus two standard deviations from the average. Statistically, this threshold will be met approximately 5% of the time during the span of measured counts. The minimum value is 3.
- **Warning Alert Threshold Hint.** This value is the average value for the counter over the specified test interval, plus three standard deviations from the average. Statistically, this threshold will be met approximately 1% of the time during the span of measured counts. The minimum value is 6.
- **Error Alert Threshold Hint.** This value is the average value for the counter over the specified test interval, plus four standard deviations from the average. Statistically, this threshold will be met approximately 0.3% of the time during the span of measured counts. The minimum value is 9.
- **Critical Alert Threshold Hint.** This value is the average value for the counter over the specified test interval, plus five standard deviations from the average. Statistically, this threshold will be met approximately .005% of the time during the span of measured counts. The minimum value is 12.

Note that the statistical percentages described above assume a perfect bell-shaped curve distribution, which is generally not the case. However, the

suggested "hint" values on this screen form a good initial guess at what the threshold for a particular alert should be initially.

The specific "hint" values change as the test interval is increased or decreased. Generally, to increase the thresholds, the test interval can be increased, or to decrease the thresholds, the test interval can be decreased. Increasing the test interval averages more messages and decreases false positives, but also decreases the time to alert.

The "Message Statistics" screen is available only for catalog items (that is, for Devices, Users, Facilities, and Severities and Threads). Clicking on a catalog item and then clicking the "View Catalog Statistics" hyperlink at the bottom of the page can also access this screen. In particular, the user can access the "Message Statistics" screen for any device as follows: The user clicks the device IP hyperlink anywhere in the system to display the device information screen; the user clicks the "View All Messages For Device" hyperlink at the top of the device information screen to show the catalog of messages for the device; the user clicks the "View Catalog Statistics" hyperlink at the bottom of the device message catalog.

Alert Auto-Learn Function

As described above, the user can access each alert on the system, view the alert threshold hints, and make adjustments accordingly in order to tune the system alert thresholds.

However, this operation can also be performed automatically using the "Alerts > Config > Auto-Learn" screen, which will automatically execute each night and make alert adjustments to alert thresholds based upon the alert severity, test interval, and message statistics.

This function is enabled by default, and will examine each alert on the system. For each alert, the auto-learn function will inspect the associated catalog of messages for the alert, derive statistics, and then adjust the threshold in order to match the alert hints. This serves to automatically tune the system to the correct alerting levels based upon historical data.

Users can disable auto-learn for the entire system, or for any specific alert item. (Alert items each have an "Auto-Learn Enable / Disable" setting incorporated into their "Alert Edit" screen.) The Auto-Learn function will suppress adjustments until the associated alert catalog has a certain number of records (by default 100 messages.) The Auto-Learn function will automatically stop once the threshold for an alert has been adjusted a certain number of times (by default after 10 days.) Any alert threshold of "1" is automatically bypassed from auto-learning, since this threshold is used to indicate a singular occurrence of a message (such as the event logs being cleared a single time.)

The Auto-Learn function is especially useful for an "unattended" role for the CorreLog system. The user can install CorreLog, direct system messages to the CorreLog platform, and then return after a period of time to see the particular alert settings and tickets that have been opened following the Auto-Learn period.

Alert Message "Suggest" Button

An additional feature of the CorreLog "Alert" function is its ability to suggest alert messages. The "AddNew" and "Edit" Alert dialogs provide a "Suggest" button that will automatically derive an appropriate alert message for the user based upon the selected counter name and type, compare function, compare value, and severity of the alert.

The user clicks the "Suggest" button, optionally edits the alert message, and then clicks "Save" or "SaveNew" to apply the new message to the system. Note that this derived message is only intended to be a valid baseline for the alert message, and the user may wish to further refine the message before saving the alert, incorporating specific corrective action or other site specific text within the alert message to indicate the precise nature of the alert condition.

Trigger, Device Alert, and Pattern Instances

Generally, CorreLog manages messages in aggregate form, without special regard to IP address or message source during correlation. For example, an alert to track invalid logins is triggered (in the absence of any special match pattern) across all messages without regard to their individual IP address. Such an alert will be fired if more than the threshold number of logins occurs irrespective of the particular device.

Experience indicates this is one of the best ways to manage messages (since it covers situations where an attacker might be attempting to break into multiple machines, for example.) However, it may sometimes be the case that alerts are to be tracked on a "per device" basis only. One way to accomplish this is to simply set up a thread and alert combination for each individual device. However, this technique may require massive numbers of alerts and threads to cover all the devices in an enterprise.

CorreLog provides the concept of "Instance" related to Triggers, Device Alerts, and Patterns as follows:

1. For those screens supporting instances, the system individually tracks messages by IP address, spawning a new instance of the trigger, device alert, or pattern when a new IP address is detected.

2. As messages are received, new instances are added, and existing instances are updated. When a threshold is reached for an instance, an alert is generated for that instance, and a ticket is opened.
3. While the alert is set, no further tickets are opened for the instance. When the alert is cleared, the instance is deleted and the above process can start over again.
4. The screens reflect the rolled-up status of the instance, and provide a link to view the active instances for the alert.

Using the above technique, the operator can create aggregate alerts, or can enable alert instances. The "Alerts > Devices" screen operates in a fashion similar to the "Alerts > Counters" screen, but supports this concept of instances. The "Alerts > Patterns" screen permits the user to set the context of the pattern detection to be either "All Messages", or "Same Device".

If a screen permits instance management, as described herein, the top-level status of the alert is displayed (and updated in real-time via AJAX), and the operator can drill down to view the instances for the alert via the "Active Instances" hyperlink that accompanies the alert. This allows the user to view the status of all instances, and acquire more information on the alert instances, from the top-level screen.

Correlation Rule Dependencies

Correlation elements are often dependent on other elements. For example, "Alerts" depend on "Threads"; "Threads" can depend upon "Triggers", and both "Threads" and "Triggers" can depend upon macros.

CorreLog prevents the user from deleting a correlation item that has dependents. For example, if a macro is used in various locations, you cannot delete the macro without first modifying those dependents. Clicking the "Edit" button for the element, and then clicking the "View Dependents" hyperlink at the bottom of the edit screen can view the particular dependents for any given correlation element.

Additional Correlation References

This section has outlined the major components and usage of the correlation feature, however substantial additional information is available in the "Advanced Correlation System User Guide", which is available in the "s-doc" directory of all Correlog installations, also accessible via the "Help" link at the upper right of the display, and at the CorreLog website.

The "Advanced Correlation Guide" document provides a discussion of many superior correlation features of the system, not documented in other locations, including detailed explanations of correlation match expressions, triggers, macros and global variables, all of which permit a higher degree of custom correlation. That manual is intended for more sophisticated users and developers wishing to implement complex correlation. Refer to that manual for higher information.

Section Summary and Additional Notes

1. CorreLog provides "semantic" correlation to detect the specific meaning of a stream of messages in real time, so that consequential real-time actions can occur.
2. The purpose of the "Correlation / Threads" component is to categorize event messages, arriving at the CorreLog Server, into "Threads" based upon simple or complex expressions
3. The purpose of the "Alerts / Counters" component is to detect when thread counters (or other counters on the system) increment above or below a certain rate.
4. The purpose of the "Correlation / Trigger" component is to detect when a particular event occurs, and then open a "gate" that enables further threading and actions.
5. The above components work together, connected by the event log, to achieve a high degree of correlation. Users connect these components together, using the Syslog as the common area of communication.
6. Users add a correlation rule by clicking the "AddNew" button. Users can view, modify or delete a correlation rule by clicking the "Edit" button at the left of the correlation rule.
7. Correlation screens all permit complex expressions, which can be simple keywords or wildcards or may be a more complex expression supporting "and", "or", "xor", "not", and containing field matches and parenthetical nesting of expressions
8. All correlation expressions are case insensitive, and permit partial matches.
9. The "Threads" and "Actions" components contain the ability to match a "Trigger" state of either "Set" or "Clear". This furnishes a special ability to "gate on" or "gate off" the correlation counters and actions based upon a

previously received event message, which establishes a context for subsequent messages.

10. The "Alerts" function permits the user to establish thresholds on system counters, which sends Syslog messages back into the system and opens tickets. This function includes "Threshold" hints and a "Suggest" function for the actual alert message (and ticket) text.
11. The user can define "Macro" expressions, which simplify the setup and maintenance of match expressions. The user creates a macro (representing a match expression) and then uses this macro in the "Match Expression" field of "Threads", "Triggers", and / or "Actions".
12. The user can define "Address Groups", which are used in "Thread" and "Actions" screens. These are similar to "Macros", but represent lists of addresses (possibly wildcarded) that are included or excluded from the match.
13. Both "Macros" and "Address Groups" use the naming convention @@name@@, where double at signs delimit the start and the end of the macro expression. However, macros cannot be used as address groups, and address groups cannot be used as macros.
14. When implementing address groups, the user should be careful not to overuse this facility. Messages can be managed by device groups, but more typically, messages are managed by their textual content.
15. The user cannot delete a correlation item that has dependents, without first removing or fixing these dependencies. The user can view the dependents of "Threads", "Triggers", "Macros", and other items by clicking the "Edit" button, and then clicking the "View Dependents" hyperlink at the bottom of the resulting edit screen.
16. Additional information on CorreLog correlation features is contained in the "Advanced Correlation System User Guide", available in the "s-doc" directory of CorreLog installations.

Section 6: CorreLog Automation

The previous section discussed correlation features of the CorreLog Server, including a brief discussion of "Actions". This section provides more detailed information on CorreLog Actions, which is the ability of the CorreLog Server to launch specific programs when certain conditions indicate.

The CorreLog Server, in addition to being a data collector and correlator, can also be viewed as an "agent" program, which runs autonomously to perform specific tasks in response to event messages. These actions can provide notifications, data processing, and corrective actions. The screens supporting this capability are available via the "Correlation > Actions" navigation tab.

In particular, this section describes one of the most flexible and important aspects of the CorreLog Server, which is how to not only process information for viewing, but also use the real-time collection of messages to perform specific actions and automation. This greatly expands the potential role of CorreLog within an organization.

This section will be of interest to administrators, wanting to configure the CorreLog Server using one of the out-of-box actions, as well as system integrators and developers that want to extend the system to perform highly specialized functions in the context of a larger management goal. Although ordinary operators may find this information interesting, they will be limited with regard to actually using the information herein, since configuring actions is limited to CorreLog users that have an "admin" login to the system.

CorreLog Actions Configuration Screens

Actions are configured via either the "Correlation > Actions" screen, or the "Ticket > Actions" screen. These two facilities are similar, but operate on different types of data: The "Correlation > Actions" screen processes incoming messages, whereas the "Ticket > Actions" screen processes tickets.

The "CO-catlog.exe" program reads the action configuration file. This process (which is also responsible for cataloging data on the system) stops the parsing of incoming messages, executes the specified action, waits for the action to complete, and then continues. All data is buffered while the CO-catlog.exe program has stopped, so no loss of messages occurs. However, if the action takes a long to execute, it is possible that the CO-catlog.exe program can begin to lag behind in its cataloging process. This situation, while unlikely, should be a consideration when adding actions to the system. In particular, actions should always be transient processes that execute within a second or two. If an action program blocks execution, the cataloging of data and further processing of actions will stop.

It is possible (and in fact common) for a single event message to result in multiple actions being executed. For example, a message may result in e-mail, the issuance of an SNMP trap, and a database update, as well as other user defined actions.

The "Add New Action" Wizard

The simplest and fastest way to configure one of the pre-existing actions is to simply click the "Wizard" button on the "Correlation > Actions" and "Ticket > Actions" screens.

The "Add New Action" wizard guides the user through the process of configuring one of the existing actions on the system. In addition to allowing the operator to select one of the pre-configured actions (such as LOGFILE, SENDMAIL, etc.) the wizard will perform checks, and provide a prompt for the specific arguments required for the action program.

This may be satisfactory in many (and perhaps most) situations. However, the CorreLog "Actions" component is an open system. It is designed to permit user specified action programs (such as Java, PHP and Perl scripts) to be executed in response to correlated items, as discussed in the sections that follow.

Detailed Steps In Configuring A User Action

It is quite common for users to employ their own actions. Hence, the majority of this section will discuss the detailed steps needed to create a brand new action.

(Note that the "Add New Action" wizard, described above, works only with the existing pre-configured actions, and not user defined actions.)

To illustrate, the following basic steps need to occur in order to add an action to either the "Correlation > Actions" or the "Ticket > Actions" screens.

1. The administrator defines the action by creating a program in the "CorreLog/actions" or "Correlog/t-actions" directories (depending upon whether this is an action on incoming messages, or an action on tickets.) This can be an existing action, or a new action, such as the execution of a "Java", "Perl" or "PHP" script, or some batch file or executable
2. The administrator copies the action into the "CorreLog/actions" or "CorreLog/t-actions" directory. The administrator might wish to test the action by executing the program at a command prompt.
3. The administrator goes to the "Actions" screen of the CorreLog web interface, and enters the match patterns that will trigger the action, the name of the action program residing in the appropriate directory, as well as any arguments for the selected action program.
4. The administrator optionally tests the action by sending a Syslog message (possibly with the "Messages > Config > Params > Send Syslog Message" tool) and verifies that the action is executed. The status of the executed action appears on the "Actions" screen, and a log file is created in the "actions" directory.

Each of the above basic steps is discussed in more detail within the paragraphs that follow.

Using An Existing Action Program

The CorreLog Server comes with various actions already prepared, described in detail within later sections of the program. Each preconfigured action is in the form of a batch file, which is commented and readily modifiable, and launches an internal program within CorreLog. Each action requires specific arguments, and receives as data certain environmental variables on the system.

The basic preconfigured actions, discussed in more detail later, are introduced below. Some of these actions are available as both correlation actions and ticket actions, whereas others are specific only to correlation or tickets, as noted below.

- **LOGFILE.bat.** This action program simply logs the triggering event message to a user specified file. It is the simplest action, and simply requires the user to provide a log file pathname on the system as an

argument. This action is available as both as a correlation action, and a ticket action.

- **RUNSQL.bat.** This action program updates a database table name MessageData with the message information. The action program requires the pathname to a configuration file as the first argument. More detailed information on using this command is provided later in this section. This action is available as both a correlation action, and a ticket action.
- **SENDLOG.bat.** This action program sends a Syslog message to another Syslog server (perhaps another copy of CorreLog running on a different platform, or a Unix syslogd process listening on the network.) The action program requires a single argument, which is the hostname or IP address of the device to send the Syslog message. This action is available as both a correlation action, and a ticket action.
- **SENDMAIL.bat.** This action program sends an e-mail message containing the message information. The action program requires two arguments: the first argument is the hostname or IP address of an SMTP server, and the second argument is the e-mail address of the recipient of the e-mail message. This action is normally available only as a ticket action.
- **SENDTRAP.bat.** This action program sends an SNMP trap message corresponding to the message information. The action program requires two arguments: The first argument is the hostname or IP address of a network manager that is running a trap receiver (listening to the standard SNMP port of 162). The second argument is the trap community name for the trap. A short MIB definition file, describing the trap, is included later in this section. This action is available as both a correlation action, and a ticket action.
- **WINMSGGR.bat.** This action program uses the "net send" command to pop a window onto the screen of a computer. The target computer must be running the standard Windows Messenger service (not to be confused with Live Messenger, or some other Messenger service.) The action program requires a single argument, which is the Net BIOS name (i.e. hostname) of the computer system that will receive the popup message. Note that this program requires the Messenger service to be started at the target computer. This action is normally available only as a ticket action.
- **RELAY.bat.** This action is similar to the "SENDLOG.bat" action described above, except relays a syslog message to another copy of CorreLog, preserving the device IP address of the message. (If the SENDLOG.bat" is used, the IP address of the message becomes the CorreLog server that is sending the message.) This action program is mainly useful in a multi-level management strategy where CorreLog installations at a local level

relay messages to a master enterprise version of CorreLog. This action is available as both a correlation action and a ticket action.

- **TUNNEL.bat.** This action is similar to the "RELAY.bat" action, described above, except uses the CorreLog Encrypted TCP Tunnel process to send messages. The program sends a TCP message to the CO-trecv.exe" program running at a remote CorreLog installation, preserving the IP address of the message. This action is mainly useful in a multi-level management strategy, where TCP is desirable rather than UDP. This action is available as both a correlation action and ticket action.
- **HELPDESK.bat.** This action can be used to interface the CorreLog ticketing system with an external helpdesk (in particular BMC Remedy, but also others.) The user must manually add the logic associated with the helpdesk incident insertion, creating or modifying a batch file called "HELPDESK_U.bat" in the "t-actions" directory. Assistance from CorreLog support may be required. This action is normally available only as a ticket action.
- **SEND_EPO.bat.** This action can be used to interface the CorreLog ticketing system with McAfee ePolicy Orchestrator. (This is a separately licensed option to CorreLog.) The configuration of this action is documented in a separate manual, included with each copy of CorreLog. (See the "More > User Manuals" hyperlink for this manual.) This action is normally available only as a ticket action.

Creating A New Action Program

The above action programs may be all that is necessary for a site. However, one extreme power, offered by the CorreLog Server, is the ability to script specialized actions. An administrator can create sophisticated and highly customized action programs to augment the above set of predefined programs. These custom action programs can be written or scripted in any language.

The starting point for adding a new action may be one of the batch files mentioned in the previous section. The "TEMPLATE.bat" file, included in the "CorreLog/actions" and "CorreLog/t-actions" directory may also be used to begin the development effort.

Correlation Actions Environmental Variables

When the program is launched via the "Correlation > Actions" program, the following environmental variables are set prior to the action program's execution, and can be used by the action program as part of its function. The list of environmental variables for "Correlation" Actions is provided below. (Refer to the section on CorreLog tickets for a similar list associated with "Ticket" actions.)

- **S_TIME.** This is the time of the message, in YYYY/MM/DD HH:MM:SS format. It is the time that the message was received (which may be slightly different from any time embedded in the message.)
- **S_TIMESTAMP.** This is an integer number representing the number of seconds elapsed since midnight (00:00:00), January 1, 1970, coordinated universal time, according to the system clock. This may be useful for certain actions that update database tables.
- **S_ADDRESS.** This is the IP Address of the device that sent the message.
- **S_FACILITY.** This is the textual description of the facility code that was incorporated within the message. The value is in lower case, and is identical to the facility names shown on the various CorreLog screens.
- **S_FACILITY_NUMBER.** This is identical to the S_FACILITY value, except is the facility expressed as an integer number in the range of 0 to 23. If the facility is a user defined facility code, then the value will be 1 (corresponding to the standard "user" facility code.) The S_FACILITY value, described above, will reflect the name of the actual user-defined specified facility.
- **S_SEVERITY.** This is the textual description of the severity code that was incorporated within the message. The value is in lower case, and is identical to the severity names shown on the various CorreLog screens.
- **S_SEVERITY_NUMBER.** This is identical to the S_SEVERITY value above, except is the severity expressed as an integer number in the range of 0 to 7.
- **S_MESSAGE.** This is the message portion of the received message, and can be 1024 characters or less.
- **S_USERNAME.** This is the user name (if any) parsed from the message. The username is a value that is contained in the "Messages > Users" tab of the program. If no user name exists in the message, the value is blank.
- **S_SUMMARY.** This is similar to the message, except that all punctuation marks are removed, spaces are collapsed, and the text is limited to 120 characters. This is mainly useful for entering text into a database table. (See the RUNSQL.bat discussion, below.)
- **S_ACTION_NAME.** This is the value of the action name, configured on the "Edit Correlation Action" screen. This can be used as meta-data by the script, independent of any other arguments to the script.

- **S_SANITIZE.** This value is either "Yes" or "No", reflecting the value selectable at the bottom of the "Edit Ticket Action" screen. If "Yes", then the value of environmental variables are "sanitized" (i.e. special characters such as "< > |" and others are removed or substituted to assist with batch file programming. If set to "No", then the message contains special characters. (See additional notes below on sanitizing environmental variables.)

The above environmental variables are set prior to the execution of any action program by the CO-Catlog.exe service, and are in addition to any ordinary environmental variables that exist (such as PATH, and COMSPEC, etc.)

Installing The Action

All action programs reside in the "CorreLog/actions" directory, for the case of "Correlation" actions, or the directory "CorreLog/t-actions", for the case of "Ticket" actions. For security reasons, to limit and control the particular programs that are automatically executed, CorreLog cannot execute an action program that resides outside of this directory. This limitation is small, and it is common practice for the administrator to create a wrapper, in the form of a Windows batch file, that executes arbitrary programs on the system.

Note that if the user copies "perl.exe" or "php.exe" or "rb.exe" into the "actions" directory, then the user can specify as an argument to the program the name of the appropriate script, as well as any additional arguments to that script. This provides substantial development room for sophisticated power users and programmers.

Adding A Correlation Action At The Web Interface

Initially, there are no actions configured for the system. The administrator clicks on the "AddNew" button to add a new action, and then fills out the "Add New" action screen to provide the action. This screen is similar to the "Correlation Threads" screen except that the user specifies both the name of a batch file or executable program in the "CorreLog/actions" directory, along with any command line arguments to the program. The following screen fields are provided:

- **Match Time Range.** As with the "Correlation Threads" screen, the user selects a time range. The event must occur within the specified time window. The time window is specified as a start time and an elapsed time (in hours.) If the elapsed time is "+24 hrs" then all messages are matched regardless of the time of day.
- **Match IP Address.** The operator selects a match pattern for the IP address of the message. This can be a fully qualified IP address, or a

wildcard, where each dot is filled with “*” to match any subnet. For example, if the user specifies “192.168.1.1”, then only that device matches the pattern. If the user specifies “10.3.*.*”, then only devices on the “10.3” class-B subnetwork match the pattern. The user can also match address groups, defined in the "Correlation > Config > Address Groups" screen.

- **Match Facility.** The operator selects a match facility, or uses the keyword “Any” to match any facility. Only those messages with the selected facility will match the pattern. The selection permits user-defined facilities, as described in the previous section.
- **Match Severity.** The operator selects a match severity, and an optional compare function. The “Any” setting matches all severities (regardless of the compare function.) The setting of EQ Error matches only those messages with Error severity. The setting of GE Critical matches only those messages of Critical, Alert, and Emergency severity.
- **Match Keyword.** The operator selects a simple or complex expression, consisting of single words, wildcards, as well as "AND", "OR", "NOT", "XOR" conjunctions, possibly parenthetically nested. The operator can also implement one or more macro values defined in the "Correlation > Config > Macros" screen.
- **Action Program Name.** This field contains the base name of the action program, residing in the “CorreLog/actions” directory, including the “.bat” keyword. The value must be a base name, and cannot be a pathname (that is, cannot contain either forward or backward slashes, or any leading dots.) Examples of valid action programs include “SENDMAIL.bat”, “RELAY.bat”, and “RUNSQL.bat”. The system will not permit the user to select a program name that does not exist in the “CorreLog/actions” directory.
- **Action Program Arguments.** This field contains zero or more arguments to the action program. This field depends upon the particular action program, and the program usage must be consulted before any arguments are specified. Each argument must be double quoted if it contains spaces. The argument cannot contain a carriage return or newline. (This delimits the list.)

Note that, rather than relying on the setting of environmental variables, the user can pass arguments as part of the command line arguments. For example, to pass the %S_MESSAGE% environmental variable to a Perl script “myscript.prl”, following some other argument, the following might be used.

```
@C:\perlscripts\myscript.prl -check "%S_MESSAGE%"
```

In the above case, the action program would either be a batch file, or might be a copy of "perl.exe" that has been installed in the "CorreLog/actions" directory by the administrator. The Perl program will receive two arguments: the first will be the "-check" argument (a constant) and the second argument will be the textual content of the message that triggered the action.

Testing The Action

Once the action is installed and configured, generating a Syslog message can test the action. The user may find the "system/sendlog.exe" utility useful in accomplishing this. Also, the "Messages > Config > Parms > Send Syslog Message" screen allows the user to send Syslog messages via the web interface.

When an action program is triggered, the "Actions" screen records the time of the last execution, and the number of times the action program has run since system startup. The "Action" screen also records the last exit status of the program. (The "exit" batch file command can set this value, depending upon whether any errors were detected within the batch file.)

Additionally, whenever an action program executes, the CO-catalog.exe program creates a log file of any output resulting from the program. This logfile is placed in the "CorreLog/actions" directory, and given the same name as the action program, except with a ".log" suffix. This log file contains both the standard output, and error output of the executed action program. This log file is rewritten each time the action program executes.

Sanitize Environmental Variables

Because "Correlation Actions" are often used with batch files, the values of environmental variables are made "Windows Batch File Safe", by substituting or removing troublesome characters from the "S_MESSAGE", and other variables. This behavior can be changed by setting "Sanitize Environmental Variables" = "No" on the "Edit Correlation Action Screen (towards the bottom of the screen.)

By default, environmental variables are sanitized. However, if the user is executing another script interpreter (such as "Python" or "Perl" or "PHP") it will not be necessary to sanitize the values. In this case, the "Sanitize Environmental Variables" setting can be set to "No". However, note that this will make it risky or impossible to use the environmental variable value anywhere within a Windows batch file.

Using Third Party Scripting Languages

The user can specify a ".php", ".prl", ".py", or other script (not just a batch file) and reference the script interpreter with a standard "shebang" character combination. In these cases, the user simply includes the standard "shebang" operator "#!" at the top of the script to launch the interpreter.

For example, if the "perl" scripting interpreter resides in the C:/Perl/bin" directory, the user can include the following as the FIRST line of the file, and this will automatically launch the perl interpreter when the action script is invoked:

```
#!C:/Perl/bin/perl.exe
```

Note that the user can specify either forward or backward slashes, and that the path (like all Windows paths) is not case-sensitive.

Detailed Action Program Usage And Debugging

The remainder of this section provides detailed notes on the configuration and usage of each of the action programs provided with the out-of-box CorreLog Server configuration. This section will be useful to any administrator attempting to configure a particular action. The following important notes apply.

- **Action Working Directory.** When an action executes, the working directory for the action is the "CorreLog/system" directory.
- **Action Permissions.** When an action executes, the user ID is the "System" ID, which is the user ID of the persistent CO-Catlog.exe program. This may affect certain aspects of the action program, especially if the administrator has modified it, or it is a new action program being added to the system.
- **Virus Protection Issues.** One of the most common problems with executing actions is that the administrator is unaware of Virus Protection programs that are executing on the system, which are blocking the action. For example, if the platform is executing the McAfee Net Shield program, the administrator must make sure that the program has permission to execute, as configured in the "Access Protection Properties" of the "Virus Scan" console.
- **Firewall Issues.** Another common problem with executing actions is that a firewall is preventing communication with the SENDLOG.bat, SENDTRAP.bat, or SENDMAIL.bat (or other) action programs. The network administrator may need to reconfigure firewalls to permit access to certain port numbers.

Experience shows that configuring action programs is quite simple, and that any problems related to the execution of these programs are generally due to one of the above considerations, or lack thereof.

The CO-Action.exe Background Program

Finally, it should be noted how actions are actually executed by the CorreLog server. Both the CO-catlog.exe and CO-action.exe programs must be executing for an action program to be launched.

The CO-catlog.exe program is responsible for queuing action programs, by creating ".bat" files on the system. These ".bat" files are subsequently launched by the CO-action.exe program. The CO-catlog.exe and CO-action.exe programs communicate through a simple file locking system that is a traditional approach to handling queues. The actual queue directory is the Correlog\queue folder, which is visible to administrators on the CorreLog Server platform.

Normally, actions are executed within milliseconds after a message is received. However, it is possible that a faulty action program (or some other environmental problem) can delay execution queued actions, causing numerous queued actions to exist. For example, if a mail server goes offline, there may be a significant delay before the mail server returns online.

To prevent a flood of actions from collecting within the queue, the "Max Queued Actions" setting of the "System > Parms" screen limits the number of action programs in the queue. If the "Max Queued Actions" value (by default 1000) is reached, all the messages in the queue are deleted, and CorreLog will log an internal error message (which appears as a regular message in the "Search" screen and other locations.) This indicates a likely problem with a user action script.

RUNSQL.BAT: Updating SQL ODBC Databases

One of the most useful actions that can be configured is to update a relational database table with received messages. This is accomplished via the RUNSQL.bat action program.

For example, a user may wish to selectively update a relational database table with any message that meets a particular criterion or pattern. (If multiple criteria exist, the RUNSQL.bat program can be configured multiple times in the list of actions, each with the same arguments, but matching different specific patterns.) With small modifications, the user can also perform database updates, or other more sophisticated SQL operations.

The RUNSQL.bat program is one of the more difficult actions to configure but is still quite simple. The steps to configure this action program are listed below:

1. The Administrator creates a System DSN, using the standard Windows > Control Panel > Administrative Tools > ODBC tools interface. The DSN must be a "System" DSN, and the name is arbitrary. An existing system DSN can be used.
2. The Administrator edits the "CorreLog/system/RunSQL.cnf" file to provide the name of the ODBC data source above, and any username and password required to access the database. (See notes below.)
3. The Administrator creates a database table called MessageData in the target relational database. This table contains four columns, "Time", "Address", "Facility", "Severity", and "Message". (See further notes below.)
4. The Administrator adds the RUNSQL.bat action to the list of CorreLog Actions, specifying the "runsql.cnf" file as the single and only argument to the program. (The other data is passed via environmental variables to the batch file.)

The RunSQLcnf file is a simple configuration file, similar to the various other configuration files within the CorreLog Server. However, the file resides in the "system" directory, hence can be specified as an argument to the RUNSQL.bat program without a full pathname. (The ".cnf" suffix is required.)

This file must be created or adjusted by the administrator, such as via "notepad.exe" or some other text file editor. A template for this file is provided in the basic CorreLog installation directory, in the CorreLog/system directory. The file contains one mandatory field (the dsname field) and two optional fields (the username and password fields, if so required.)

An example of this configuration file is provided below. This file contains the following fields.

```
dsnname      CorreLog-database
username     none
password     none
```

In the above example, the ODBC Data Source Name, created via the Windows Control Panel in step #1, is "CorreLog-database". This value was identified when the "System DSN was created. In this case, the username and password are both "none" (indicating that a login is not required for the database that the "CorreLog-database" DSN represents.) The ODBC data source "CorreLog-database" is maintained by the Windows system, and contains the actual location of the database, as well as any other parameters that are required to access the database, such as logins, and network locations.

The MessageData table, which is required and updated by the RUNSQL command, is created as follows. (Note that this is preprogrammed into the RUNSQL.bat action program, and can easily be tailored and modified by a user to change the table name, column types, column widths, and other information.)

```
create table MessageData
(
    MessageTime          varchar(20) ,
    MessageAddress       varchar(16) ,
    MessageFacility      varchar(12) ,
    MessageSeverity      varchar(12) ,
    MessageData          varchar(128)
)
```

The above format is somewhat arbitrary. More efficient types of tables can be created, but for most applications (of perhaps 1 million messages or less) the above table definition should be adequate to contain the pertinent messages matched by the Action facility. NOTE that the message is truncated to 128 characters automatically.

Finally, note that RUNSQL.bat action program, although quite simple, can be expanded upon to include multiple DSN files (each with a different configuration file argument) and performing various types of updates, such as updating counter values, creating table joins, and deleting old data on the system.

SENDLOG.BAT: Relaying Messages

It may be the case that the CorreLog Server is not the central collection point for messages, and merely participates in a larger management strategy, where various Syslog and Trap Receiver agents roll information up to higher levels, filtering and processing this information along the way.

In this case, the SENDLOG.bat file can be used to relay the received message to another Syslog server. This program can be configured as a program action, possibly matching a select number of messages based upon severity, facility, hostname, time of day, or keyword.

The SENDLOG.bat program takes the specified Syslog message and then sends the message to another server. The Syslog message is sent with the same severity and facility as the original message. The receiving host will see the message coming from the CorreLog platform.

To configure the SENDLOG.bat action program, simply add the program with a single argument, which is the destination of the Syslog message. This is sufficient to begin relaying Syslog messages.

Note that if the system receiving the Syslog message is “localhost”, then the message will immediately trigger a second invocation of the action program. This will cause the de-duplication function of the program to kick in. If the user modifies the “SENDLOG.bat” program to prefix the message with the name of the sending host, this can cause problems because each received message is different than the last, and the event log will be flooded with messages, at least temporarily until the system stabilizes with a duplicate message.

The SENDLOG.bat program makes use of the “CorreLog/system/sendlog.exe” program, which is a useful command line utility for sending Syslog messages, and which has brief help built into the program. This program takes four arguments: the destination host, the message, the optional severity, and the optional facility. (The severity and facility must be represented as numbers.) If no severity or facility is specified, they default to “debug” and “user”, respectively.

Execute the sendlog.exe program at a command prompt, with no arguments, to see how the program is used. Also, try using the sendlog.exe program interactively to send Syslog messages to the CorreLog Server. (This is a valuable exercise, and illustrates how easily Syslog messages can be sent by user written scripts and other programs.)

SENDMAIL.BAT: Sending E-Mail Messages

It may be that certain events are so critical that an operator wishes to be notified of the event via e-mail or a pager. The SENDMAIL.bat action program can accomplish this. When configured as an action, the program will connect to an SMTP server, and send an e-mail message to a single configured user. The default message is suitable for use with an alphanumeric pager, or for collection by a mail program. To send to multiple users, the administrator can configure additional instances of the SENDMAIL.bat program, or can modify the batch file to include CC or BCC directives.

The SENDMAIL.bat program requires three arguments: The first argument is the hostname or IP address of an SMTP server, the second argument is the e-mail address of the person sending the e-mail message, and the third argument is the recipient of the e-mail message. NOTE that the second argument (the person sending the e-mail message) IS REQUIRED for operating most SMTP servers, and MUST be a valid e-mail address, such as the address of the CorreLog administrator.

The content of the e-mail message, by default, consists of a subject (which indicates the severity) followed by the body of the message (which indicates the time, the hostname, the facility, and the message.) This is configurable in the SENDMAIL.bat program. An administrator can modify the content of the message, including the Subject line.

In particular, an administrator can copy the SENDMAIL.bat program to the SENDPAGE.bat file, hardcode the SMTP server and / or e-mail address, and possibly omit the subject line to make the message appear better on a text pager. The SENDMAIL.bat program is well commented, and designed for easy adaptation.

The SENDMAIL.bat program makes use of the "CorreLog/system/sendmail.exe" program, which is a useful command line utility for sending e-mail messages, and which has brief help built into the program. This sendmail.exe program takes four arguments: the destination host, a "from" e-mail address a "to" e-mail address, and the pathname to a file that contains the message to send. Within the file that is sent, the standard e-mail headers of "Subject:" "CC:", "BCC:" and other items can be placed prior to composing the message.

Given that, the operation of the SENDMAIL.bat program is easily understood. SENDMAIL.bat simply constructs a temporary file, executes the sendmail.exe program, and then deletes the temporary file. (Because the CO-Catlog.exe program processes each event sequentially, there is no possibility of one device conflicting with another device, and the temporary filename does not have to be unique.)

Note that anti-spam filters almost always require a valid “From” address. Also note that it is common for virus protection programs to interfere with email programs. (See notes at the start of this section.)

The typical reasons for e-mail messages failing to be sent are as follows:

1. The “From” e-mail address may need to be set to the name of a valid user on the system. Some SMTP servers will deny access to the system unless both the hostname is known, and the “From” address indicates a valid user of the system. In that case, the administrator should make sure that the second argument to the SENDMAIL.bat program is a valid e-mail user.
2. A virus protection program may be blocking the transmission of e-mail messages. All virus protection programs have an “Exclude” setting that permits the user to enter the “sendmail.exe” program as a valid program for sending e-mail on the system. Either disable the virus protection program, or add the “sendmail.exe” program as exclusion to the virus checker.
3. The destination user is not reachable, or is invalid, or the SMTP server is not valid. This is a surprisingly common problem that is easily fixed by verifying the correct SMTP server name, the “from” e-mail address, and “to” e-mail address are specified as the three arguments to the SENDMAIL.bat program.

One easy way to test communication to the SMTP server is to use “telnet” at a command line prompt, and connect to port 25 of the SMTP server. For example, the user can issue the command “telnet smtp-server 25” and see if the telnet session can be started. This simple technique can be used to troubleshoot a large measure of problems.

Another way to test the program: The user can interactively execute the “CorreLog/system/sendmail.exe” program at a command prompt, and see what error messages (if any) result. These error messages will also appear in the “actions/sendmail.log” program whenever the SENDMAIL.bat action program is executed.

Finally, note that there are no safeguards to prevent a user from flooding an e-mail server or pager with messages. It is the important responsibility of the user to configure the program action with enough restrictions so that ONLY pertinent messages are sent. This can typically be accomplished by using this action program ONLY with a Severity setting of GE-Critical, which substantially reduces the amount of e-mail messages that will be generated.

SENDTRAP.BAT: Sending SNMP Trap Messages

One important way for the CorreLog Server to interoperate with other network managers on the network is to send SNMP trap messages to those managers. This is accomplished with the SENDTRAP.bat action program.

The SENDTRAP.bat action, when configured on the “Actions” screen, will send a standard SNMP V2 trap to a trap receiver on a remote (or the local) platform whenever a message is received that matches the configured pattern. The user can elect to send all messages as SNMP traps, or a range of messages. The user may wish to configure multiple instances of SENDTRAP.bat actions to match complex patterns based upon devices, text strings, severities, or other criteria.

The SNMP Network Manager trap receiver should be listening for SNMP traps at the standard port number of 162. The trap message, sent by the CorreLog Server, contains a variable binding with the message attached, and the trap number is the severity number of the message that initiated the trap, ranging from 0 to 7. This should make the reception of SNMP traps, and the decoding of these messages, somewhat trivial.

To configure the SENDTRAP.bat action program, simply add the program with two arguments: the first argument is the trap destination (i.e. the hostname or IP address of the platform running the trap receiver.) The second argument is the trap community. The trap community is a required argument to the SENDTRAP.bat program, even if the manager does not require it. (The industry standard trap community, which will generally work for most managers, is “public”.) In general the trap community be a short text string, with no spaces, that serves as a password to the trap receiver. (If you have questions, contact the administrator of the SNMP network manager, who will give you the correct value.)

The SENDTRAP.bat program makes use of the “net-snmp/snmpttrap.exe” program. This is a standard utility of the Net-SNMP open source software. The location of this software is specified in the “Messages > Config > Parns” screen, and is by default the directory “CorreLog/net-snmp”. Although the SENDTRAP.bat program is configured to send SNMPv2 traps, it can also be reconfigured to send SNMPv3 traps, or Inform type traps. This information is well documented within the standard Net SNMP distribution software.

If the SNMP network manager requires MIB definitions for the trap, the user can compile the official trap definitions at the SNMP network manager. These trap definitions reside in the “CorreLog-doc/CorreLog-MIB.txt” file in standard ASN.1 format. As shown in this file, the enterprise object used in the trap is “experimental.514”, and the trap contains a single variable binding of “sysDescr.0”, which contains the originating message hostname and text.

TUNNEL.bat: Relay Message Via TCP Tunnel Process

In a multi-tier management strategy, it may be usefully for one CorreLog Server to send messages to another CorreLog server. This can be accomplished via the "SENDLOG.bat" or "RELAY.bat" actions, or can be accomplished via the "TUNNEL.bat" action described here.

Each of these techniques is similar (in that syslog messages are sent to another copy of CorreLog) but slight differences exist. The TUNNEL.bat file operates in a fashion almost identical to the RELAY.bat program, accept requires the CO-trecv.exe program (CorreLog Tunnel Receiver program) to be executing at the remote location. The message is then direct to the tunnel receiver, and sent via TCP rather than UDP, guaranteeing reliable delivery

To configure the TUNNEL.bat action program, simply add the program with a single argument, which is the destination of the Tunnel process. This is sufficient to begin relaying Syslog messages to another CorreLog Server site.

Note that if the system receiving the Syslog message is "localhost", then the message will immediately trigger a second invocation of the action program. This will cause the de-duplication function of the program to kick in. If the user modifies the "TUNNEL.bat" program to prefix the message with the name of the sending host, this can cause problems because each received message is different than the last, and the event log will be flooded with messages, at least temporarily until the system stabilizes with a duplicate message.

The TUNNEL.bat program makes use of the "CorreLog/system/tunnel.exe" program, which is a useful command line utility for sending messages to CorreLog via TCP, and which has brief help built into the program. This program takes four arguments: the destination host, the message, the optional severity, and the optional facility. (The severity and facility must be represented as numbers.) If no severity or facility is specified, they default to "debug" and "user", respectively.

Execute the tunnel.exe program at a command prompt, with no arguments, to see how the program is used. Also, try using the tunnel.exe program interactively to send Syslog messages to the CorreLog Server. (This is a valuable exercise, and illustrates how easily Syslog messages can be sent by user written scripts and other programs.)

Administrators can easily modify the TUNNEL.bat to affect the particular action. This batch file is mainly included for completeness, as a simple example of a general-purpose action, as opposed to an action that is useful in the general case. Given that, some sites may find this action program quite useful.

TICKET.bat: Automatically Opening Tickets

Normally, CorreLog tickets are opened exclusively by the "Alerts" screen in response to either message rates, or patterns, or some other criteria. This provides assurance that tickets are opened only on exception, and multiple tickets are not opened until the exception is cleared. For example, while some persistent condition (such as a network scan) is occurring, the alert remains in a "Set" condition, and no further tickets are issued until the persistent condition ceases.

In some situations, it may be advantageous to open tickets directly when a message is received, so that the ticket is purely event driven. For example, if the user wants to open a ticket on every occurrence of a message, it is difficult to configure an alert to accomplish this. (The "test interval" for an alert may be set small, such as one second, but can not be set to zero seconds.) To open a ticket when a particular message is received, the operator can implement the "TICKET.bat" file as a Correlation Action, as described herein.

To configure the TICKET.bat action program, simply add the program with three arguments: the first argument is the person assigned the ticket, the second argument is the textual content of the ticket, and the third argument is the numeric severity of the ticket ranging from 0=emergency to 7=debug. When a particular message is received, the TICKET.bat file opens the ticket, uses the message as the "Related" message, and runs any configured ticket actions. The ticket textual content accepts the \$T_ variables used by other alert programs, so the related message, related device, related device description, and other values can be incorporated into the content of the ticket.

The TICKET.bat program makes use of the "CorreLog/system/OpenTicket.exe" program, which is a useful command line utility for opening tickets at a command line. (See Section 7 for additional notes on this program.) The "OpenTicket.exe" program has additional application in automation, such as in the creation of more complex correlation alerts that directly open tickets. In this case, the user can copy the TICKET.bat file and make adjustments that may be necessary to accomplish a specific correlation activity, or create a new action script that directly executes "OpenTicket.exe".

To prevent generation of multiple tickets by a script or program, the "TICKET.bat" file relies on the "Ticket Parameters" configured by the operator. The operator can adjust the value of "Max OpenTicket.exe per minute" (by default 20 executions per minute) up or down to block a runaway situation that might otherwise generate excessive numbers of tickets.

Section Summary And Additional Notes

1. Two types of actions exist: "Correlation" actions and "Ticket" actions. These types of actions operate in a similar fashion, but accept slightly different information and are launched differently.
2. Correlation Actions are configured via the "Correlation > Actions" screen, and are executed when specific messages are matched. Correlation Action programs reside in the "CorreLog/actions" directory
3. Ticket Actions are configured via the "Ticket > Actions" screen, and are executed when specific tickets are opened, closed, and / or changed. Ticket Action programs reside in the "CorreLog/t-actions" directory.
4. Action programs generally consist of batch file wrappers for other programs, accepting information and environmental variables from the CorreLog system, and launching other programs.
5. The user can update a relational database table with message information using the RUNSQL.bat action program. This action requires the user to configure an ODBC data source, and also requires the user to configure the "system/RUNSQL.cnf" file. The RUNSQL program requires a single argument, which is the pathname to the RUNSQL.cnf file.
6. The user can send log messages to another Syslog server using the SENDLOG.bat file. This program requires a single argument, which is the hostname or IP address of the Syslog host.
7. The user can send e-mail messages using the SENDMAIL.bat. This program requires three arguments, the SMTP server hostname or IP address, the "Mail From" e-mail address, and the "Mail To" e-mail address.
8. The user can send SNMP traps using the SENDTRAP.bat file. This program requires two arguments, the SNMP manager hostname or IP address, and the trap community for the trap.
9. The user can send messages via the Windows Messenger, creating popup alerts on the local or a remote machine. This program requires a single argument, which is the Net BIOS name of the platform that is to receive the alert message. The target platform must be running the standard "Messenger" service.

Section 7: CorreLog Tickets

The Correlation functions, discussed in the previous section, organize the data into related groups and patterns. The raw message list is correlated into distinct threads and alerts.

This correlated data, while important, may not necessarily be "actionable". That is, the data may be of interest to auditors and useful in forensic studies, but the data does not necessarily suggest a problem that requires manual intervention by a specific person.

The CorreLog "Tickets" function can be thought of as the highest level of correlation, where specific correlated patterns generate incident tickets that are assigned to specific users and groups. In many cases, these tickets represent the final output of the CorreLog system, performing the final step of assigning a problem or incident to a specific user or department for analysis for corrective action. As such, one final item is correlated with the message data, that item being an incident "assignee".

The CorreLog Ticket function, discussed here, is similar to traditional ticket systems. Incidents are automatically opened by the system, and can be manually closed when a resolution for the incident is found.

Additionally, as discussed herein, the CorreLog Ticket function can be interfaced to other ticketing systems to operate in a larger incident management strategy. CorreLog interoperates with a variety of third-party systems, and common management databases, using a simple script interface detailed within this section.

Basic Ticketing Concepts

It is possible for a ticket to correspond to a single targeted message. But more commonly, tickets correspond to multiple types of messages, possibly associated with various threads. The operator should not confuse a ticket with a specific message, but rather a specific pattern of messages that have occurred during a particular time interval. This greatly expands the ability of the Ticketing system to identify specific threats and incidents.

Ticket generation is a direct function of the "Alert" facility, discussed previously. Alerts monitor various counters on the system, especially "Thread" counters". Alerts then generate Syslog messages back into the system when thresholds are violated. Additionally, when a Syslog message is generated, the user can optionally assign the Syslog message to an individual or group. This is configured on the "Alert Editor" screen, shown below.

The screenshot displays the 'Alert Editor' screen in the CorreLog Server web interface. The browser window title is 'CorreLog Server - Microsoft Internet Explorer'. The interface includes a navigation bar with tabs: Home, Dashboards, Messages+, Correlation+, Alerts+ (selected), Tickets+, Reports+, and System+. Below this is a sub-navigation bar with tabs: Counters (selected), Devices, Patterns, Custom, and Config+. The main content area contains the following fields and controls:

- System Counter Name:** A dropdown menu currently showing 'Unselected'.
- Compare Function:** A dropdown menu showing '(GE) Greater Than Or Equal'.
- Threshold:** A text input field containing '3' with the label 'Counts Per Interval'.
- Test Interval:** A text input field containing '60' with the label 'Seconds'.
- Match Alert Time:** A dropdown menu showing 'Midnight' and a '+ 24 hrs' button.
- When Threshold Is Triggered...** (Label only, no input field visible).
- Send Alert Message:** A text area with a 'Suggest...' button and a '256 characters available' indicator.
- Insert Alert Variable:** A dropdown menu showing 'None' and an 'Insert...' button.
- Enable Auto-Learning:** A dropdown menu showing 'enabled'.
- Alert Facility:** A dropdown menu showing 'alert'.
- Alert Severity:** A dropdown menu showing 'notice'.
- Assign Incident To User:** A dropdown menu showing 'admin'.

At the bottom of the interface, there is a 'Done' button and a 'Local intranet' status indicator.

The above screen is accessed by clicking the "Correlation" tab, then the "Alerts" tab, and then the "Edit" (or "AddNew") button. The user selects a counter name, a compare function, a test interval, and specifies an alert message, facility and severity.

To cause the alert to additionally open a ticket, the operator configures the "Assign Incident To User" field to be the name of the ticket recipient. This causes the ticket (with the specified alert message) to be created for the user and to appear in the "Tickets" tab of the program as an open ticket. By default, the ticket assignee is the name of the user who created the alert, but the value can be any registered user of the CorreLog system (configured in the "System > Logins" tab of the program.) The Ticket Assignee can also be a group, defined in the "Tickets > Config > Ticket Groups" screen.

Not all alerts will necessarily open tickets. The operator can configure the value of "Assign Incident To User" to be "disabled", in which case no ticket is actually opened. For example, it may be desirable or necessary to have a series of alerts (feeding messages back into the Syslog) to be further correlated, and a ticket opened only if the correlation yields a certain count. In this case, various alerts may be disabled to prevent tickets from being opened due to intermediate correlation steps.

Ticket Assignee

The Ticket assignee receives the alert message, which appears on the "Tickets" tab as an open ticket. If the user has set the "Initial Ticket Group" in his or her user preferences, the ticket is displayed at the top of the list, and only those tickets assigned to the user are displayed. (The operator can view other ticket groups via a pull down menu on the "Tickets" screen, including the special "All" group.)

Any CorreLog user may be assigned a ticket. Additionally, assignees can be created via the "Tickets > Config > Ticket Groups" screen. For example, the tickets may be related to an operational group such as "Routers", or "Web Servers", where users with particular expertise in this group can collect and view ticket information.

The Ticket Assignee can be modified, and the ticket assigned to some other user. The operator edits the ticket via the "Update #" button, changes the ticket assignee, and then saves the ticket. The ticket will be moved from the current ticket group to the new group, causing the ticket entry to disappear from the current "Assigned To:" group view.

Ticket Groups

When creating an alert of any type, the "Assigned To" field specifies who the ticket is assigned. This is a drop-down menu on all the "Alert" screens.

The "Assigned To" field can be any CorreLog defined user. However, the "Assigned To" field can also be a functional unit in the organization, such as "DB_Admin". Ticket groups are useful several different ways.

- **Organizing Tickets.** The user can add and use a ticket group to assist with the organization and "tagging" of ticket data with a username. Specifically, the operator can view the tickets assigned to any user on the system (of which he or she may have an interest.) The Assigned To field appears as a drop-down item on the "Tickets > Opened" and "Tickets > Closed" fields.
- **Organizing Ticket Actions.** When a ticket action is created (via the "Ticket > Actions" screen) the action can match the "Assigned To" group. This permits a user to receive e-mail notifications for a particular group, or other actions that include the group name to be executed. The Ticket Group appears as an environmental variable in all ticket actions scripts, facilitating custom scripting.
- **Tracking Status.** The number of opened and closed tickets, and total ticket count is available via the "View Groups" link at the top and upper-right corner of the "Tickets > Opened" and "Tickets > Closed" screen. The operator can click on this link (in a fashion similar to the "View Groups" link on "Devices" and "Threads" screens) to see the relative activity associated with each ticket group.

Creating good ticket groups can facilitate the general organization of the system, and assist with complex notifications associated with sending e-mail from certain types of tickets, or other notifications. Experience shows that creation and maintenance of good "Ticket Groups" are well worth the administrative effort.

Editing Tickets

The user can edit a ticket to change the assignee, the ticket message, the severity, and to supply a text message resolution. The user can also set the status of the ticket to either "Opened" or "Closed". This permits the user to record pertinent facts about the ticket, such as what corrective action was necessary.

The "Resolution" field of the ticket can be set to any arbitrary text string to explain how the situation was resolved. This value is permanently connected to the ticket, but not checked for accuracy. The value may be important when using

third-party incident management systems, as discussed in later sections. The default resolution value is "None."

Closing And Deleting Tickets

When tickets are closed, they are removed from the "Opened" tab, and added to the "Closed" tab. The closed ticket contains all the information of the open ticket, and can be further edited, such as to change the ticket resolution, or to change the ticket status back to "Opened".

To close a ticket, the user clicks the "Update #" button for the ticket, then sets the ticket status to "Closed" and clicks the "Update" button. As an alternative, the user can close all the tickets (currently displayed on the top-level screen) using the "Close All" button beneath the list of opened tickets. This provides a fast method of closing all tickets, and generates an "Audit" message to the Syslog indicating the username and source IP address of the user that closed these tickets.

Note that any closed ticket can be re-opened by clicking on the "Closed" tab, clicking the "Edit #" button for the ticket, and then setting the ticket status back to "Opened".

Finally, administrators can delete a ticket rather than close the ticket. This removes the ticket permanently from the system rather than moving the ticket to the "Closed" list. This function is available only to administrative users.

Manually Adding Tickets

Normally, tickets are opened automatically by the system by the "Alert" facility. However, It may be convenient or necessary to manually open a ticket, such as to test ticket automation, or annotate a problem that is visible in the system but has no corresponding correlation logic.

The user can manually open a ticket by clicking the "AddNew" button, which permits the user to add arbitrary text, assign the ticket to a particular user or ticket group, and assign a severity to the ticket. In this case, the ticket is identical to a ticket that has been automatically opened, except that the ticket has no related messages or source alert definition (described below.)

Related Messages And Source Alert Definition

Associated with each ticket (except for tickets that are manually added, as described previously) are one or more messages that triggered a particular alert on the system.

The user can view the messages that triggered the ticket by clicking the "Related Messages" hyperlink on the main screen (also available on the "Update" screen.) This displays the messages responsible for opening the ticket, including the device, the time of the message, and message content. These messages are permanently affixed to the ticket and cannot be changed.

The user can view the source alert that opened the ticket by clicking the "Source Alert Def" hyperlink. This displays the alert counter, threshold, and other parameters. If the user is an administrator, the user can directly edit this alert configuration from the "Source Alert Def" screen, and these changes will apply to any future tickets opened by the system.

Note that the related messages can (and often do) span various devices and message contents depending upon the alert and counter definitions that have been configured.

Further note that only the messages directly related to the ticket's opening are displayed, and not necessarily the latest messages from the device. The number of related messages is exactly equal to the threshold of the source alert definition. To view the latest messages, the user can click the "Related Messages" hyperlink on the top-level "Tickets" screen, then click the device IP address hyperlink, and then click the "All Messages For Device" hyperlink on the resulting device information screen.

Ticket Actions Facility

The power of the CorreLog ticketing interface is greatly expanded by its ability to interface directly to third-party incident management systems. This is accomplished via the "Ticket > Actions" facility, which is executed each time that a ticket is automatically or manually opened, and each time that the ticket is subsequently modified.

The "Ticket > Actions" facility is similar to the "Correlation > Actions" facility except that the environmental variables instantiated for the batch file are slightly different, reflective of the current ticket information. The specific environmental variables are enumerated below.

- **T_UID.** This is a unique identifier for the ticket. The value may be useful in interfacing with a third-party management system. The value is permanently associated with each particular ticket instance, and never changes.
- **T_STATUS.** This is the keyword "Opened" or "Closed" or "Changed", indicating the current ticket operation by the system or user. The value of "Changed" indicates the ticket has been changed (but not necessarily closed.)

- **T_STATE.** This is the keyword "Opened" or "Closed" indicating the current ticket state. Note that the T_STATUS value can be "Changed" and the current state can be either "Opened" or "Closed" (if the user modifies a closed ticket.)
- **T_DATE_TIME.** This is the time when the ticket was first opened, in YYYY/MM/DD HH:MM:SS format. This can generally be taken as the time of the last message, which caused the ticket to initially open.
- **T_TIME_NOW.** This is the time of the action's execution, in YYYY/MM/DD HH:MM:SS format. This will be the same as the T_DATE_TIME value when the ticket is opened, and will reflect the time of any additional change or ticket closure.
- **T_TSTAMP.** This is similar to the T_TIME_NOW value, but is the time of the action's execution represented as a long integer value, which may be of use in certain applications such as updating a relational database with information.
- **T_ASSIGNED_TO.** This is a string indicating the current assignee for the ticket, as displayed on the top-level screen.
- **T_SEVERITY.** This is the severity of the ticket in text format. The user assigns the severity in the alert definition, but can later change the severity to be any valid value.
- **T_SEVERITY_NUMBER.** This is similar to the T_SEVERITY value, but is the severity of the ticket in numeric format, where 0=emergency, and 7=debug.
- **T_MESSAGE.** This is the ticket message, identical to the message displayed by CorreLog on the "Tickets" screen.
- **T_SUMMARY.** This is similar to the value of T_MESSAGE, except that all punctuation marks are removed, spaces are collapsed, and the text is limited to 120 characters. This is mainly useful for entering text into a database table.
- **T_COMMENT.** This is any text associated with the "Resolution" field of the ticket, under 256 characters, identical to the resolution text displayed by CorreLog on the "Tickets" screen.
- **T_EXTURL.** This is the URL to the CorreLog server that opened the ticket (including "http:" prefix.) This value is useful for incorporating into e-mail messages, or otherwise referencing the particular CorreLog server.

- **T_RELATED_ADDRESS.** This is the address of the device that caused the ticket to be initially opened. This value is available only when the ticket is first opened. It will be the address of the most recent message in the "Related Messages" for the ticket.
- **T_RELATED_DEVNAME.** This is the device name corresponding to the T_RELATED_ADDRESS value (above), defined in the Device Information screen. If no device name is defined, this value is a zero length string.
- **T_RELATED_DEVDESC.** This is the device description for the device corresponding to the T_RELATED_ADDRESS value (above), defined in the Device Information screen. If no device description is provided, this value is a zero length string.
- **T_RELATED_FACILITY.** This is the textual facility of the message that caused the ticket to be initially opened. This value is available only when the message is first opened. It will be the textual facility of the most recent message in the "Related Messages" for the ticket.
- **T_RELATED_SEVERITY.** This is the textual severity of the message that caused the ticket to be initially opened. This value is available only when the message is first opened. It will be the textual severity of the most recent message in the "Related Messages" for the ticket.
- **T_RELATED_USERNAME.** This is a user name, parsed from the message that caused the ticket to be initially opened. This value will be provided ONLY if the related message (below) contains a username that appears in the "Messages > Users" tab, otherwise there is no value.
- **T_RELATED_MESSAGE.** This is the content of the message that caused the ticket to be initially opened. This value is available only when the message is first opened. It will be the content of the most recent message in the "Related Messages" for the ticket.
- **T_RELATED_SUMMARY.** This is similar to the value of T_RELATED_MESSAGE, except that all punctuation marks are removed, spaces are collapsed, and the text is limited to 120 characters. This is mainly useful for entering text into a database table.
- **T_ACTION_NAME.** This is the value of the action name, configured on the "Edit Ticket Action" screen. This can be used as meta-data by the script, independent of any other arguments to the script.

- **T_SANITIZE.** This value is either "Yes" or "No", reflecting the value selectable at the bottom of the "Edit Ticket Action" screen. If "Yes", then the value of environmental variables are "sanitized" (i.e. special characters such as "< > |" and others are removed or substituted to assist with batch file programming. If set to "No", then the message contains special characters. (See additional notes on sanitizing environmental variables.)

A template for ticket actions, TEMPLATE.bat is provided on the system within the CorreLog "t-actions" directory. The batch file can be modified to launch external programs (such as the ARS Remedy "RunMacro.exe" program, or ASG ASF.exe program for working with the ASG IMPACT system.) Note that the user can immediately execute any program within the batch file, such as "Perl" or "Java", to take precise action on the ticket information.

Finally, note that it is easy to test the operation of a ticket action. Simple add a ticket to the system via the "AddNew" button. This will cause the ticket action file to be executed, including the passage of proper data to the program. The results of the ticket program's execution (if any) will be available via the "View Debug Log File" hyperlink for the action, available on the "Ticket > Actions" screen.

Ticket Group Wizard

The Ticket Group Wizard, accessed via the "Tickets > Config > Ticket Groups > Add New" button, provides a simple way to create a simple correlation rule and ticket group consisting of a single correlation thread and alert. This may be adequate for many applications.

Note that the process for creating a ticket is as follows:

1. The user optionally creates a ticket group for the ticket, perhaps identifying a particular fault.
2. The user creates a correlation thread and targets specific messages via a match expression or macro.
3. The user creates one or more alerts, which monitor the counts for the correlation thread, and opens the ticket for the user.
4. The user optionally repeats steps 2 and 3 to configure multiple correlation threads, and multiple alerts for each thread counter.

The above steps are not difficult. However, it is often the case that a single ticket group corresponds to a single thread and alert. (This is especially true when first getting started with the CorreLog system.) In this case, the user can create the ticket group, thread, and alert values using the "Ticket Group Wizard" screen in one easy session.

The wizard permits the user to either create a new correlation rule, or modify an existing correlation rule (referenced by a ticket group.) The wizard cannot configure sophisticated correlation and ticketing strategies, which may be necessary or practical in some enterprises. However, using the wizard, the operator can quickly configure multiple ticket groups, and later modify or enhance these rules as necessary.

Auto-Close Duplicate Tickets Setting

In many situations, conditions cause multiple tickets with the exact same content to be opened on the system. This is a normal situation, and the user will typically handle this situation by addressing each individual ticket as a separate incident, fixing this problem at the source. The Auto-Learn function, discussed in a previous section, will work to reduce the number of tickets opened on the system by adjusting the thresholds of alerts up or down, thereby increasing the thresholds for those alerts that are commonly opening tickets.

If the user wishes to modify this behavior, so that each ticket opened on the system is unique, the user can enable the "Auto-Close Duplicate Tickets" setting, found on the "Tickets > Config > Parms" screen. When this setting is enabled, the system operates as follows:

1. When a ticket is automatically opened, all the open tickets for the same day as the ticket are automatically inspected by the system.
2. If an earlier ticket with exactly the same content is found, that earlier ticket is closed when the new ticket is opened.
3. The closed ticket has a resolution automatically added, indicating that the ticket is a duplicate, and has been closed.

The effect of this operation is that all the tickets on the "Open" screen are unique, and the most recently opened ticket is at the top of the list. A history of tickets can be found in the "Closed" tab of the system.

The following special caveats apply:

1. Only tickets that are automatically opened by the system are de-duplicated. Duplicate manual tickets can be opened on the system (via the "AddNew" button on the "Ticket" screen).
2. When a ticket is automatically closed by the system due to a duplicate ticket, any action associated with the closed state is not executed. This reduces the amount of e-mail on the system, and other actions.
3. The auto-closed tickets must have the same day as the new ticket. The Auto-Close Duplicate Tickets function does not span more than one day,

and only operates on the current day. Duplicate tickets can exist for different days, but only one duplicate ticket per day.

4. The exact same content for the ticket is required. If the user has inserted a variable (such as the related message) into the alert message, the entire ticket content must precisely match. In particular, if the user inserts the ticket UID as part of the alert message, that ticket will never be duplicated (because the UID is always unique to the ticket.)
5. Ticket actions are often executed when a ticket is opened, hence the number of opened tickets may not correspond to the number of e-mails on the system.
6. No check is made of the related messages for a ticket. Only the main ticket content is checked. Hence, the related messages for an automatically closed ticket may be different than the currently opened ticket. This may require inspection of closed tickets on the system.

Sanitize Environmental Variables

Because "Ticket Actions" are often used with batch files, the values of environmental variables are made "Windows Batch File Safe", by substituting or removing troublesome characters from the "T_MESSAGE", and other variables. This behavior can be changed by setting "Sanitize Environmental Variables" = "No" on the "Edit Ticket Action Screen (towards the bottom of the screen.)

By default, environmental variables are sanitized. However, if the user is executing another script interpreter (such as "Python" or "Perl" or "PHP") it will not be necessary to sanitize the values. In this case, the "Sanitize Environmental Variables" setting can be set to "No". However, note that this will make it risky or impossible to use the environmental variable value anywhere within a Windows batch file.

Searching for Ticket Information

Tickets are distinct from other messages. Although each ticket generates a message when it is opened (via the "Alert" facility) the related message and assignee of the ticket are not generally added to the message logs, hence the "Message Search" function may not be adequate to completely qualify a particular ticket.

There are two main facilities available for searching for tickets.

1. The "Ticket Advanced Search" screen is accessed via the "Tickets > Opened" screen, and permits the user to search for a ticket by UID, Assignee, related message, and other attributes of the ticket. This facility allows the user to identify the ticket by various values of interest.
2. The "Query" screen provides a "Ticket search" capability that allows the operator to search for tickets like other messages. Additionally, the operator can make use of the various forensic aspects of the "Query" screen, such as the "Analyze" and "Graph" functions.

The "Ticket search" capability of the "Query" screen operates on the "Ticket Archive" files, which includes all the tickets that have been archived (instead of all the tickets opened recently on the system.) These archives are the textual form of the ticket, updated at midnight each night, and kept for "Keep Archive" days (as configured on the "Messages > Config > Parms" screen.)

Common Management Database Setup

As an advanced function of the Ticket facility, the user can configure the tickets to automatically update a Common Management Database (CMDB) where the status of each ticket group is automatically updated in an SQL database at periodic intervals, by default within 10 seconds after any ticket change. The CMDB data provides the following fields.

- **Enable CMDB.** This field must be set to "Yes", and the CorreLog system must then be restarted. (This will schedule the CO-CMDB.exe service to be executed, which is required to run the CMDB function.)
- **Data Source Name.** This field must be set to the name of a configured DSN that provides permissions to the CorreLog program. Any ODBC compliant database may be used.
- **Database Username.** This field must be set to the database username, matching the value configured in the DSN above. If no username exists, the value should be set to "None"

- **Use Database.** This field must be set to the database segment name, matching the value configured in the DSN above. If no database segment name exists, the value should be set to "None".
- **Database Update Seconds.** This field is the time, in seconds that controls when the database is updated after a change. By default, the value is 10 seconds, indicating that the database is updated within ten seconds after any change. Setting this value too low can unnecessarily consume CPU resources.
- **Current Count Interval.** This time specifies when a ticket is regarded as "Current", by default one Hour.
- **Hour Count Interval.** This time specifies when a ticket is regarded as "Recent", by default four Hours.
- **Severity Display Mode.** This determines whether the worst-case ticket group severity is expressed as a percentage ranging from 100 (best) to 0 (worst), or a standard severity in the range 0=emergency to 7=debug.

CMDB Table Definition

Once the CMDB has been enabled, and CorreLog system restarted, the CO-CMDB.exe program will automatically create a database table in the target DSN, and update this database with information when ticket status changes. The program creates the following table automatically:

```
create table Tickets
(
    Username                varchar(40) ,
    Description              varchar(40) ,
    CurrentCount             integer,
    HourCount                integer,
    DayCount                integer,
    CurrentSeverity          integer,
    HourSeverity             integer,
    DaySeverity              integer,
    PercentLoad              integer
)
```

The various table columns are described below.

- **Username.** This is the name of the user that is assigned the ticket. Users are created by the System > Logins screen, and specified in the Alerts >

Counters screen. One special user is the "All" user, which does not need to be added to the system.

- **Description.** This is the full name of the user that is specified by "Username", configured on the System > Logins screen, and can be used as a label on any CMDB dashboard.
- **CurrentCount.** This is the total number of tickets opened for the user since the "CurrentCount" interval, by default 1 hour.
- **HourCount.** This is the total number of tickets opened for the user since the "HourCount" interval, by default 4 hours.
- **DayCount.** This is the total number of tickets opened for the user since midnight of the previous day.
- **CurrentSeverity.** This is the numeric severity of the worst-case ticket opened during the "CurrentCount" interval. The value ranges from 100% to 0% or 0 to 7, depending upon the setting of the "Severity Display Mode" field, discussed previously.
- **HourSeverity.** This is the numeric severity of the worst-case ticket opened during the "HourCount" interval. The value ranges from 100% to 0% or 0 to 7, depending upon the setting of the "Severity Display Mode" field, discussed previously.
- **DaySeverity.** This is the numeric severity of the worst-case ticket opened since midnight of the previous day. The value ranges from 100% to 0% or 0 to 7, depending upon the setting of the "Severity Display Mode" field, discussed previously.
- **PercentLoad.** This is a percentage of tickets assigned to this user, with respect all tickets on the system. (The "All" user will have a value of 100%.)

The above values permit a dashboard to be implemented to view the CMDB data, suitable for use by operations personnel. In particular, the table architecture permits indicator lights to reflect the status of each user group, reflecting current ticket status, recent ticket status, and daily ticket status. To affect the values, the user simply logs into CorreLog and closes tickets or changes ticket severities.

Note that proper setup of this function generally requires CorreLog professional services. Interested users should contact support@correlog.com for both instructions and the software necessary to configure the CMDB interface.

Ticket System Access Group

Administrators can give users a "ticket" type system access, which permits the user to access just the ticketing system and no other part of CorreLog. This is configurable for a user in the "System > Logins" screen (available only to "admin" type users.)

The "ticket" system access group is useful for those sites where operators are assigned tickets, and are tasked with resolving and / or dispatching tickets.

In this situation, the administrator can create one or more CorreLog users with a "ticket" type login. These users will see only the "Ticket" tab, and will be able to see the various components of the ticketing system (such as the source alert definition and related messages) but will otherwise not be able to access any CorreLog data or other screens.

Tickets External URL

The "Tickets" component has a unique URL that access ticket data without any other access to the system and, with extra configuration, without a CorreLog login. This allows users to view tickets that may be assigned to them or their group without compromising the security of other data on the CorreLog system.

The "Tickets" screen is the only CorreLog application with this capability. The feature permits limited drill-down into the CorreLog system by users who otherwise have no access to the system.

The URL to access the ticket information is as follows:

```
http://(Hostname)/x-cgi/web.exe?(Assignee)&(Listmode)
```

Note that the URL uses "x-cgi", (not "s-cgi"), which corresponds to the "x-cgi" directory of the CorreLog root directory. This special URL has three required arguments:

- **Hostname.** This is the name of the computer system that is running the CorreLog application. This is the standard hostname component of the URL, either an official hostname or an IP address that is executing the CorreLog server process. If the CorreLog web server is running at some port other than the standard port 80, the value should contain the port specification, such as "myhost:88"
- **Assignee:** This is the user or ticket group that is assigned tickets. This must be one of the assignee values of the system, either a CorreLog user name or a ticket group name. The Tickets screen will display all the tickets assigned to the specified user or group.

- **Listmode:** This is the list mode of the tickets, which appears at the top of the Tickets screen. The value is either "current", "hour-n", "lastday", or "max-n". The Tickets screen will limit the display of tickets to the specified time range or count.

For example, if the CorreLog web server is running on "myhost" at port 88, a user can display all the tickets for the "operations" group that have been opened for the last four hours using the following URL:

```
http://myhost:88/x-cgi/web.exe?operations&hour-4
```

By default, when accessing the above URL, the user must still log into the CorreLog server using HTTP authentication. (The web browser will prompt for a username and password). This required HTTP login may be disabled by an Administrator by renaming or deleting the "x-cgi/htaccess.txt" file from the CorreLog root directory. This will permit the external URL to be used by anyone without special authentication. For example, this will permit the URL to be employed on corporate intranets so that ordinary users can see the current tickets.

More information on configuring CorreLog to access special screens can be found in the "CorreLog Sigma Web Framework" Users Manual.

Tickets and Auto-Learning Function

The "Auto-Learn" function, discussed in the Section 6, has specific components that work with the ticketing function. Specifically, the Auto-Learn function can adjust alert thresholds based upon the number of open tickets on the system, raising thresholds automatically to reduce the number of tickets.

Additionally, if a ticket is opened when an alert threshold is automatically adjusted (because the Auto-Learn function has determined that the alert threshold is too low or too high based upon historical message counts) the Auto-Learn function can automatically close the ticket.

If a ticket is automatically closed by the Auto-Learn function, the Auto-Learn program will automatically create a ticket resolution, which appears as part of the closed ticket. This resolution message will indicate that the ticket was automatically closed, and will report the previous and current alert thresholds associated with the ticket.

OpenTicket.exe Command Line Executable

If the user wishes to open a ticket via a script or at a command line prompt, the "OpenTicket.exe" program can be used. This program resides in the "system"

directory of the CorreLog distribution, and is used by the "TICKET.bat" correlation action. The script must be executed in the "system" directory (in order to locate the "../ticket" folder). Otherwise, there are no special requirements to execute this program.

For example, to open a ticket at a command prompt, change working directories to the CorreLog "system" directory, and execute the following command:

```
OpenTicket.exe admin "This is the ticket text" 4.
```

The above command will cause a new ticket, assigned to "admin" with a severity of a warning (4), to be created in the "Tickets" tab. The first argument to the OpenTicket.exe program is the assigned user, the second argument is the ticket content, and the third argument is the severity of the ticket ranging from 0=emergency to 7=debug.

The operator can execute the "OpenTicket.exe" program with no arguments to receive brief help. If the "OpenTicket.exe" program is executed as a correlation action program (such as via the "TICKET.bat" action program discussed in the previous section) the related message will be added to the ticket. If the program is executed at a command prompt (without benefit of the CorreLog environment) no related message will be associated with the ticket.

Finally, because it is possible for an errant script to flood the system with tickets, the "OpenTicket.exe" program makes use of a throttle value, which limits the number of executions per minute. This throttle value is found on the "Tickets > Config > Parms" screen, and by default has a value of max 20 executions per minute.

Ticket Map Gadget

As a final topic of discussion, CorreLog provides a flexible "Ticket Map" gadget that can be used to organize tickets in a visual depiction of the ticket location. The "Map-Ticket-Gadget" program (available as a standard CorreLog dashboard gadget) allows the operator to configure an arbitrary background image (such as a geographical map) and then place icons at arbitrary locations on the image map. Each icon corresponds to a ticket group, or some number of tickets therein during a period of time.

The Ticket Map Gadget can be generalized to display the highest level of correlation (i.e. Tickets) in an arbitrary framework that gives the user abilities to support arbitrary data visualization. For example, by looking at the dashboard, the operator can determine what problem areas may exist in the enterprise, within a subsystem, or across the geographic network.

Section Summary And Additional Notes

1. Tickets provide the highest level of correlation on the system, associating a particular incident to an assignee for analysis.
2. Tickets are opened by the "Alert" facility, where an assignee for the ticket is specified. Not all alerts will open tickets, but every ticket is opened by one alert.
3. Tickets can be assigned to any CorreLog registered user. Additionally, Ticket assignees can be specified in the "Ticket Groups" screen.
4. Each ticket has a list of related messages (which caused the ticket to be opened.) The user can view the list of related messages, or the source alert definition, by clicking hyperlinks supplied on the top-level ticket screen
5. The CorreLog ticket system can interface to third-party systems via the "Ticket > Actions" facility.
6. Users can configure regular actions (such as sending e-mail) when tickets are opened, since the alert facility sends a Syslog message whenever tickets are opened.
7. The Ticket Wizard, available via the "AddNew" button on the "Ticket Groups" screen, permits the user to quickly add a correlation rule and ticket group to the system.
8. The Auto-Learn function works with the Ticketing system, automatically adjusting alert thresholds based upon the number of opened tickets, and automatically closing tickets when thresholds are automatically adjusted.
9. The "OpenTicket.exe" program can be used at a command prompt, or within a script, to open a new ticket within CorreLog.
10. The "Max OpenTicket.exe per second" setting, located on the "Tickets > Config > Parms" screen, allows the user to safeguard against rapid executions of the program that might otherwise flood the ticket log with unwanted tickets.
11. The "Map Ticket Gadget", available from the CorreLog Dashboard, allows an operator configure a graphic depiction of a geography, or schematic drawing, and then drag the state if ticket groups onto the display. This provides a mechanism for visualizing information based upon ticket states.

Section 8: CorreLog Reporting

CorreLog includes a comprehensive reporting facility, which will automatically create reports that summarize message data. Various types of reports exist, ranging from ad hoc reporting by the "Query" function, Audit reports suitable for review by auditors and analysts, and Microsoft Excel reports that automatically populate spreadsheets with data.

Reports operate in a different fashion from correlation items as follows: Reports are generally launched as background processes once each day at midnight (or on demand.) The reports document the state of the message data and system at the time the report was generated. Although similar to the correlation items discussed previously, reports are different in that they are not real-time and instantaneous. The reports refer to the statistics that existed at the specific time the report was generated.

Report screens provide "Job Control", that is the reports run in background, and once a report is launched the operator can leave the reporting screen and check back at a later time to see the results. The operator can terminate a running report, and can obtain status on the report progress.

As a special option, CorreLog will publish reports using RSS (Real Simple Syndication) protocol. An end-user can configure an RSS News Aggregators, such as the "FeedDemon" aggregator, to connect to the CorreLog server, and automatically download and archive new reports. These RSS aggregators greatly expand the ability to collect and use report data, including the ability to configure watches, and notifications when new reports arrive.

Reports Query Function

The "Reports > Query" function is the most basic type of report, and operates in a fashion similar to the regular "Messages > Search" screen. The results of the report are displayed as a list of messages that satisfy the simple or complex search criteria. The user accesses the "Run Search Query" screen by clicking the "Query" link in the upper right of the CorreLog display, or by clicking "Reports > Query" and then clicking "Run Report".

To generate a query, the operator selects match criteria on the "Run Search Query" screen, and executes the query background process to perform the search (by clicking the "Confirm" button on that screen.) The background process executes to completion (which can be within a few seconds to several minutes or more.) When the query is finished, the top-level provides results of the search, which can subsequently be analyzed using several built-in tools, including a "Graph Message" and "Analyze" function.

Report queries are typically ad hoc operations, but queries can be saved by naming the query on the "Run Search Query" screen. If the operator provides a name for the query on the "Run Search Query" screen, the search criteria are saved under the specified name. The operator can access saved queries via the "Saved Queries" link on the top level "Reports > Query" screen.

In general, a query can be used instead of a basic search (i.e. as an alternative to the "Messages > Search" screen) with identical results. The main difference between the basic search and the query is the complexity of the query, and the willingness to spend a little time in acquiring the query results. The query function provides a trade off between "speed of results" and "complexity of query". Otherwise, the two types of searches operate in a similar manner.

Detailed notes on the "Query" screen exist in the "CorreLog Screen Reference" manual, available from the "Home" screen of the CorreLog Server (available after user logon.)

Audit Reports

The "Reports > Audit" function is a simple but powerful audit facility that generates reports based upon several commonly required specifications, in particular specifications related to security compliance standards such as PCI-DSS and HIPAA (but also others.)

Audit reports generate at midnight each night, or on demand. The report generator parses raw message data to create reports that can subsequently be download in HTML format, CSV format, or text format via a "Report Viewer". Administrators publish audit reports using RSS, and can e-mail audit reports to end users.

Audit reports come "ready-to-use" and usually require no configuration by an operator. However, several configuration items exist that expand the role of audit reports to include more specialized reporting appropriate to specific administrators (that may not want or need access to all data on the system.) Operators can create alternate "Report Viewers" (via the "AddNew" button on each screen) which limit the data to certain columnar match patterns, and can hide columns of a report.

For example, the operator can create a "Database User Lockout" report that is a small subset of the "User Activity" report, limiting the report to only those database users that have a lockout condition. Likewise, an operator can create a "Critical Unix Events" report from the "Device Activity" audit reports, which lists only those UNIX platforms experiencing a certain number of critical and higher errors.

As an important side-function of "Audit Reports", the report information can be loaded automatically into ODBC compliant databases for use with third-party report writers, or to support complex SQL queries. The user configures the ODBC functions of each report type via the "Advanced" button found at the top of the various Audit reports.

Detailed notes on "Audit Report" parameters and usage exist in the "CorreLog Screen Reference Manual", available on the "Home" screen of the CorreLog Server (available after user login.) Additional application notes on the "Advanced" parameters of audit reports are available from CorreLog support on request.

Audit Report Types

The following Audit reports are accessed via the "Reports > Audit" facility of the system. Each report type is intended to satisfy a particular common audit requirement supporting good security for the enterprise.

- **User Activity.** This type of report lists all users of the system, demonstrating (in compliance with most security standards) that all users are being monitored. The report directly supports PCI-DSS and HIPAA specifications to track user access, but also provides a useful indication of user activity and possible insider threat.
- **Device Activity.** This type of report lists all the managed devices of the system, demonstrating (in compliance with most security standards) that all critical systems are being monitored. The report directly supports PCI-DSS and HIPAA specifications to track access to managed platforms. The report also provides a good indication of the amount of logging and activity associated with certain devices.

- **Perimeter.** This type of report lists all external IP addresses of the enterprises detected by router or firewall messages. The report collates all messages and their counts that contain two IP addresses (at least one of which is external to the enterprise.) This unique report tabulates the country code associated with IP addresses, useful for determining external attacks or data exfiltration.
- **Account Management.** This type of report lists all the changes detected by Active Directory with regard to adding, modifying and deleting users, and adding, modifying and deleting user groups. The report is mainly intended to support Microsoft AD implementations, but may have application in other LDAP applications (where LDAP is used to authenticate enterprise users.) This report demonstrates compliance with most security standards that account changes are monitored and managed.
- **Tickets.** This type of report lists all tickets opened and closed on the system, demonstrating (in compliance with most security standards) that threats are being managed and reviewed. The report also provides a useful indication of ticket activity on the system, which is a high-level summary of CorreLog's effectiveness in monitoring threats and security events.
- **Score Cards.** This type of report is mainly useful in demonstrating compliance to a security standard such as PCI-DSS, or demonstrating good security practices. The report maps thread counts (which represent the amount of data tabulated by the system) with respect to user selected threads and data categories. In particular, the report demonstrated data collection and organization, which is the most basic function required for any compliance standard or any good security management. As such, this type report is often the starting point for any security audit.

In particular the "Score Cards" report may be sufficient to simplify audit operations for certain types of regulatory compliance standards, since these types of reports illustrate the basic areas of coverage for the enterprise, and can be used to map specific CorreLog configuration elements to certain compliance standard specifications such as PCI-DSS, HIPAA, SOX, etc.

Excel Reports

The "Reports > Excel" function is a powerful and highly extensible reporting facility that loads Microsoft Excel spreadsheets with message statistics configured by the operator. The facility will also load small numbers of message samples into the system (but this facility is not intended to include large amounts of message data, but only message statistics.)

Excel reports generate at midnight each night, or on demand. The report generator parses raw message data to create reports and populate an Excel "template" file, to allow the user to create graphs and formulas on the various statistics residing in the message.

Excel reports operate on All Messages, but typically the user restricts the Excel report to a particular "Thread", such as "Successful Logons", or "Hardware Events", or some other thread created under the "Correlation > Threads" screen. The report then reflects statistics on a particular class of messages (rather than all messages.)

Configuring New Excel Reports, Procedure

CorreLog comes with a small number of generic Excel reports that may be sufficient for many organizations. However, it is easy to expand the number and types of reports, as follows:

1. The user clicks on the "Reports" tab to access the Excel reporting screen. This displays the top level "Excel" screen, which lists any configured reports. (The user can access these reports by clicking the report name hyperlink.)
2. On the "Excel" screen, the user clicks "AddNew" to define a new report. This displays the "Add New Report" screen.
3. On the "Add New Report" screen, the user provides a title and data source for the report. (The data sources are selectable via a pull-down menu, and correspond to each configured correlation thread. The default data source of "All Messages" can also be selected.)
4. The user saves the report configuration by clicking the "Commit" button. The newly created report will then appear in the top-level "Excel" screen.

With no other action, CorreLog will begin generating reports each midnight. The "Add New" screen provides reasonable defaults and error checking of user values. When the configuration is saved, the title for the report (specified in step #3 above) is displayed as a hyperlink, and the user can click this link to access the Excel report.

If the user does not want to wait until midnight for the new report, the user can generate an immediate report by clicking on the "GenRpt" button at the top of the "Excel" screen. This launches the report generator as a background process. The status of the background process is displayed at the top of the "Excel" screen. When the report generator finishes (which might take several minutes) all the reports have been updated with the latest information.

Excel Report Configuration Items And Settings

To create a new report, the user clicks the "AddNew" button. To modify an existing report, the user clicks the "Edit" button, next to the report entry. To delete an existing report, the user clicks the "Edit" button, and then the "Delete" button on the resulting screen. The following configuration items are provided, when the user clicks "AddNew" or "Edit":

- **Report Data Source.** This is the location that contains the data to report on. By default, the data source is "All Messages", but the user can also specify any existing correlation thread on the system. If the user selects a

correlation thread as the data source (possibly further qualifying the data with the "Match Expression" value, discussed below) only the data in that thread will be reported on. When a user creates a new correlation thread in the "Correlation Threads" screen, the name will appear in this list. If the user later deletes this thread, the value becomes "INVALID", and the report is no longer generated.

- **Report Title.** This is the title for the report that appears on the main screen, and in any configured RSS feed. The title should be under 72 characters, cannot contain any markup code or other nonstandard characters, and should be descriptive of the report content. Frequently, this value is the name of the Correlation Thread, or possibly some variation of that name. For best results, the first 20 characters should be unique (although this is not required.)
- **Startpoint Template / DSN.** This is the name of a file residing in the "excel/templates" directory. It is an Excel file that is used as the startpoint for the report. The file can contain graph objects, custom colors, formulas, pivot tables, image files, and any other custom information. By default, the system comes with a startpoint template file called "Default.xls". The user can copy this default file to a new filename in the "excel/templates" directory, and use that new file as the startpoint template for the report. Note that the startpoint template value can also be an ODBC DSN value configured in the ODBC tab, which causes report information to be loaded into a relational database rather than a spreadsheet. (See later section on support for ODBC compliant databases.)
- **Unique Table Prefix.** This is an optional three-character identifier that is used only if the Excel Startpoint Template is an ODBC DSN value. The value is accepted but not used if the "Excel Startpoint Template" value references an Excel spreadsheet. This value can be used to make table names unique when using a database as the target of the reporting function. (See later section on support for ODBC compliant databases.)
- **Match Expression.** This value can further qualify and select the data in the selected data source. The value can be a keyword, wildcard, and / or complex expression, such as "X and (Y and not Z)" as discussed elsewhere. The default value is "*" which matches all messages in the selected data source.
- **Report Span Days.** This value specifies how many days of report data are processed. The default value is 30 days. Note that, although the report is generated each night, the report contains summary data that spans multiple days. By default, a new 30-day report is generated each night, starting with the previous day. The user can change this value to span up to 500 days (if that much data actually exists on the system.) The longer

the span of days, the bigger the report file and the longer it takes to generate the report. Note that, the number of days worth of data is also limited by the "Keep Data" setting in the "Messages Config Params" screen. If the "Report Span Days" setting is longer than the actual number of days worth of data residing on the system, then the report contains data only for the existing number of days.

- **Max Data Records.** This value specifies how many records are actually processed, and can be used to limit the size of the report file. By default, the system processes up to one million messages within the specified data source. This value can be increased to a higher value, but can result in huge reports that are difficult to download or open. Experience shows that most reports contain much less than one million records.
- **Keyword Counting Rules.** This button provides access to a screen that contains arbitrary keywords, which are counted by the program. Up to four different keywords can be specified. The keyword counts are recorded in the "Keyword History" tab of the Excel spreadsheet, useful for recording specific keywords within the catalog of messages. For example, this feature can be used to look for specific errors or file types that may be of special interest.
- **Include Messages In Report.** This value specifies whether the actual raw message data is included in the report. The default value is "No", indicating that messages are not added to the report, and the report contains only message counts, device counts, username counts, facility counts, and severity counts. Setting the value to "Yes" will greatly increase the time to generate the report, and the report size. This may or may not be a consideration to the end-user. Also, including message data in publicly available reports may compromise system and network security (in some situations.) The decision to include messages in reports will be made largely on the type of data source configured and reported on.
- **Message Parsing Rules.** This button provides access to a screen that allows the user to specify parsing rules that break down messages into specific keywords and phrases. This function is available ONLY if the "Include Messages in Report" setting (above) is set to "Yes". (As described above, special cautions are associated with including messages in reports.) The user can specify up to four different parsing rules. The parsed keywords are updated in the "Messages" tab of the Excel spreadsheet. (See Message Parsing Rules section below for more information.)
- **Publish Report via RSS.** This value specifies whether the report will be published via RSS. The default value is "Yes", indicating that the report is published. For a report to be published, the user must also configure the

"RSS" parameters of the system, discussed elsewhere. This setting is mainly useful in preventing certain reports from being automatically published to RSS readers and news aggregators.

Excel Report Message Parsing Rules

On the Excel Report "Add New" screen, If the operator selects "Include Messages" to be "Yes", then the operator can access the "Message Parsing Rules" screen of the Excel report by clicking the "Advanced" button next to this option.

The message parsing rules allow the user to parse parts of a message from the stream and update the Excel spreadsheet with this data. Up to four different parsing rules can be configured for a report. Several different techniques are provisioned for parsing messages. As follows:

- **Match Patterns.** If the message is easily matched (by means of a prefix, suffix, or both) then the operator simply specifies the match pattern in the context of the message, identifying the target text to match with an asterisk (*). For example: "User Name: *" returns the text immediately following the first occurrence of "User Name:" in the message, delimited by a blank space (or other user specified delimiter.)
- **Numeric Field Positions.** If the message is not easily matched using the above technique, the user can specify a numeric field position for the target text. For example, to return the 12th word in the message, the user simply specifies the number 12 as the match specification. Words are delimited by one or more blank spaces. The actual match word is delimited by a blank space (or other user specified delimiter.)
- **Special Parsing Functions.** If a numeric field is used, the user can "cast" the value to return just a specific portion of the matched phrase. The following special functions can be used. (Each function is preceded by a '\$' character.) Importantly, the user can match a range of message fields with the special \$message() function, as documented below.

Special Message Parsing Functions

Several special functions can be used to parse messages, and can be configured with the message parsing rules. Each of these functions is preceded by a dollar sign "\$" character, followed by a numeric field position. The syntax of the special message parsing functions is as follows:

\$message (fieldpos, count)

This parse function returns all words in the message from the specified starting position, continuing for "count" words or the end of the message. If "count" is not specified, then all words to the end of the message are returned.

\$match (keyword, parse-spec)

This function requires two arguments: the first argument specifies a match keyword or phrase; the second argument is a parse specification. The function returns the results of the parse specification ONLY if "keyphrase" appears somewhere in the message content. Otherwise, the parse specification is not matched. For example, "match (apache, 3, server: *)" finds the word following "server:", where both "apache" and "server:" both occur in the message.

\$skip (keyword, count, parse-spec)

This parse function requires three arguments: the first argument specifies a match keyword or phrase; the second argument specifies the number of times to skip that keyword in a message; the third argument is a parse specification. The function returns the results of the parse specification following the "count" occurrence of the "keyword" value. For example, "\$skip (User, 3, Account Name: *)" finds the word following "Account Name:", where "Account Name:" must first be preceded by three occurrences of the word "Test".

\$integer (fieldpos)

The integer portion of the word given by "fieldpos" is returned. If the "fieldpos" value does not return a word, the match specification is not matched.

\$alpha (fieldpos)

The ASCII text portion of the word given by "fieldpos" is returned. If the "fieldpos" value does not return an ASCII value the match specification is not matched.

\$basename (fieldpos)

The portion of the word given by "fieldpos" is returned, except for the last portion of the word delimited by a punctuation mark. For example, if the matched word is "README.txt", then only the word "README" is returned.

\$suffix (fieldpos)

The suffix portion of the word given by "fieldpos" is returned, delimited by a punctuation mark. For example, if the matched word is "README.txt", then only the word "txt" is returned.

\$url (fieldpos)

The URL portion of the word given by "fieldpos" is returned. The URL must begin with "http://", "https://", or "ftp://". If the "fieldpos" value does not point to a URL, then the match specification is not matched. This is useful for matching and parsing a URL from the input message

\$dom (fieldpos)

Similar to the \$url function above, except just domain name portion of the URL portion of the word given by "fieldpos" is returned. This function returns the text following "http://" (or other URL indication) delimited by the first slash, if any. This is useful for matching and parsing the target host contained in a URL, such as when monitoring HTTP proxy servers.

\$ipaddr (fieldpos)

The IPv4 portion of the word given by "fieldpos" is returned. For example, if the matched word is "192.168.1.100:55", then only the value of "192.168.1.100" is returned. If the "fieldpos" value does not return an IP address then the match specification is not matched.

\$email (fieldpos)

The e-mail address portion of the word given by "fieldpos" is returned. (The e-mail address must be in the form "X@Y.Z".) If the "fieldpos" value does not return an e-mail address the value the match specification is not matched.

\$geo (fieldpos)

The geo-location (two letter country code) for the IP address given by "fieldpos" is returned. For example, if the matched word is "145.1.1.1", then the value of "CA" (the country code for Canada) is returned. If the "fieldpos" value does not return an IP address then the match specification is not matched.

\$delim (delimspec, fieldpos)

The portion of the message given by "fieldpos" is returned, where fields are delimited by either "space", "colon", "semi", "dash", "sp-dash", or "punct". (The value of "sp-dash" delimits fields by a space, followed by a dash character; the value of "punct" delimits fields by any punctuation.)

Note that the "Excel" reporting function is not the only CorreLog application that makes use of the above parsing functions. These functions also have application in specific areas such as dashboard gadgets, and in the "Query" reporting tool. Also, note that these functions should not obscure the relative simplicity of the parsing system. Most parsing rules may consist of a single number (rather than a complex function) such as simply specifying "4" to return the fourth word of the message.

Custom Reports Using Excel Startpoint Templates

CorreLog comes with a ready to run report file that plots Message rates, summary counts of devices, usernames, facilities, and severities, and historical counts. This report does not need to be configured or modified. However, the user can add additional pre-configured reports, or modify the existing template.

When CorreLog generates a report, it starts with an Excel file residing in the "excel/templates" directory. The file is, by default, the "Default.xls" file. The user can create and select other default start files via the "Excel" Add / Modify screen dialogs. These other start files can contain additional graphs, formulas, and formatting.

The reporting mechanism uses Excel "Named Areas" to update data. These areas are added to the Excel spreadsheet, within the Microsoft Excel tool. (The user outlines an area in the spreadsheet, then using the Excel "Insert > Name > Define" menu item.) Each area is populated by the reporting facility, and any graphs or pivot tables that use the fields within this area are automatically updated when the user views the Excel spreadsheet.

The user can reposition these named areas, but cannot rename or create new named fields. The characteristics of these fields are similar to regular SQL database tables. The following named areas are required in all Excel Startpoint template files:

- **ReportTitle.** This "named area" contains the name of the report, as displayed on the web interface screen. This "named area" must contain a single column. The "named area" typically resides on the "Home" tab of the worksheet.
- **ReportTime.** This "named area" contains the date and time that the report was generated, and also contains the system contact (as configured on the "System Params" tab of the web interface.) The area must contain a single column. The "named area" typically resides on the "Home" tab of the worksheet.
- **DeviceSummary.** This "named area" contains a list of all devices in the filtered data source and a count of the messages that have been received from the device. The "named area" must contain three columns. This typically resides on the "Dev Summary" tab of the worksheet.
- **UserSummary.** This "named area" contains a list of all the users in the filtered data source and a count of the messages that have been received that reference each user. The "named area" must contain two columns. This typically resides on the "User Summary" tab of the worksheet.

- **FacilitySummary.** This "named area" contains a list of all facilities in the filtered data source and a count of the messages that have been received from the facility. The "named area" must contain two columns. This typically resides on the "Facility Summary" tab of the worksheet.
- **SeveritySummary.** This "named area" contains a list of all severities in the filtered data source and a count of the messages that have been received with the severity. The "named area" must contain two columns. This typically resides on the "Severity Summary" tab of the worksheet.
- **FacilityData.** This "named area" contains a list of all facilities in the filtered data source, and a count of the messages that have been received from the facility, broken down by day. The "named area" must contain twenty-seven columns. This typically resides on the "Facility Hist" tab of the worksheet.
- **SeverityData.** This "named area" contains a list of all severities in the filtered data source and a count of the messages that have been received with the severity, broken down by day. The "named area" must contain ten columns. This typically resides on the "Severity Hist" tab of the worksheet.
- **KeywordData.** This "named area" contains a list of the keyword counts for the messages, broken down by day. (Up to four different keywords can be tracked, configured via the "Advanced" button on the Excel edit screen.) The keyword counts are broken down by day. The "named area" must contain five columns. This typically resides on the "Keyword Hist" tab of the worksheet.
- **MessageData.** This "named area" contains the selected Message Data, consisting of the Message Time, Message Address, Message Facility, Message Severity and Message Data fields. The "named area" must contain six columns. This typically appears on the "Messages" tab of the worksheet. This is populated in the worksheet ONLY if the report has the "Include Messages In Report" setting to "Yes". (The default is "No", to conserve processing time and file sizes.)

Each of the above Excel named areas must exist, and must have the correct number of columns. Otherwise, the Excel report will not be generated and the particular error indicating the problem will be displayed when the user attempts to access the Excel report.

Further Notes On Excel Spreadsheets

If you are only casually familiar with Excel, you should know that it contains many sophisticated statistical and analytical functions, graphing capabilities, and

formatting options. For example, the user can create formulas and functions, and create many different types of graphs. The user can also create "pivot" tables that permit a high degree of flexibility in selecting fields and re-arranging columns and rows. These Excel features can be configured into an Excel startpoint template, and will be available when the user views the spreadsheet report.

For advanced programmers, Excel incorporates a sophisticated macro and programming interface that uses Microsoft Visual Basic to create complex functions and depictions.

An abundance of information exists on Excel usage, including many published books, and a variety of website sources. For more information, consult these resources.

Manually Running Report Generators

Once the user has configured an Audit, Excel, or E-Mail report, it is automatically generated by the system shortly after midnight each and every day, and will contain report data for the previous day (or other "Report Span Days" setting.) When the report generator executes, it overwrites any existing report. A maximum of one report is retained in the "reports" directory. (The user can configure the RSS facility, discussed later, to collect multiple copies of the report, for historical purposes, in the "rss" directory folder.)

The report generation process can take a few seconds to finish, a few minutes, or under some circumstances many minutes to execute. The execution time depends upon the number of reports, the number of messages in the data sources, the configured "Report Span Days" setting for each report, and whether raw message information is included in each report. As the report generation process runs, the status of the background process is displayed at the top of the reporting screen.

If the user wants to run the report generator immediately, and not wait for midnight, the user can click on the "Run Report" button at the top of the report screen. This launches the report generator immediately. In this specific case, the date of the reports is the time that the user manually ran the report generator (rather than for the previous day) and the report will incorporate only data up to that time. Note that this manually generated report will be overwritten at midnight, when the scheduled report executes. The final scheduled report will contain the complete information for the day.

If the report generator is already running, and the user wishes to terminate the process, the user can click on the "Run Report" button (as if to launch the process) and the system will prompt to terminate the background process. This provides complete manual control over the report generator process.

If an error terminates the report generation process prematurely, the last error message is displayed at the top of the report screen, along with the date and time of the error. If an error is encountered when processing a single report (such as a corrupt Excel Startpoint template) then the report specific error is available by clicking on the report title hyperlink.

RSS Report Distribution System

In addition to generating reports, the CorreLog system includes a method of distributing (or "publishing") these reports to end-users. The CorreLog system uses RSS (Really Simple Syndication), which is a standards-based, XML oriented publishing scheme that works with a variety of clients.

RSS creates what is referred to as a "feed" or "channel", which can be read using software called an "RSS reader," or "RSS aggregator." Using this client, the user subscribes to a feed by entering the feed's URL link into the client. The client RSS reader then checks the user's subscribed feeds regularly, and downloads any new documents to the end-user.

A variety of information exists on RSS, available through simple web searches. A number of free and highly capable RSS readers are also available, including the "FeedDemon" reader. (Check the CorreLog website, under the "Resources" tab for links to this reader.) Using RSS, an end user simply needs to know the home page of the CorreLog system to begin accessing published reports.

The publishing process appears similar to the report generation process, but performs an entirely different purpose. The RSS publishing software, on a daily, weekly, or monthly basis, copies the latest Excel reports from the private "reports" directory to the public "rss" directory, providing a unique URL for each report. The publisher then configures the RSS feed XML documents to point to these new reports, permitting users to access these reports with or without a login to the CorreLog server.

The system can retain multiple copies of reports for historical analysis. Unlike the Excel and Audit reporting facility, which retains only the latest report information, the RSS facility can save up to 500 copies of each report, creating an archive of report files. Even if there are no subscribers to the RSS data, this single function may justify its configuration by the CorreLog administrator.

RSS Configuration Items And Settings

To configure RSS, the user navigates to the "Correlation > Reports > RSS" screen, and clicks the "Edit" button. The following RSS configuration items should be inspected or adjusted by the user:

- **RSS Master Feed Enable.** Default Disabled. This setting controls the RSS publishing, and must be set to "Enable" for any publishing to occur. By default, the value is "Disabled", which means that the administrator must enable RSS publishing after installation before any data is published. Note that, in addition to this setting, publishing can be controlled for individual documents using the "Publish Via RSS" setting for each individual report.
- **RSS Feed Name.** Max 40 characters. This text string is used as the official name of the feed, which appears on the RSS reader. This can be any arbitrary string, including the site or organization name. The value cannot contain any special characters or HTML.
- **RSS Feed Description.** Max 128 characters. This text string is used as the official description of the feed, which can appear on the RSS reader and other locations. This can be any arbitrary string, including the site and organization and contact. The value cannot contain any special characters or HTML.
- **RSS Item Prefix.** Max 40 characters. This string is used as a prefix to all feed items, and can be used to trigger special "watches" or "clippings" associated with the RSS feed reader. It is also used as the "Category" value in the RSS reader. The value cannot contain any special characters or HTML.
- **RSS Server Name.** Server hostname and port number. This is the name or IP address of the CorreLog server HTTP server, and can include an optional port number (such as "myserver:88"). The value should be a permanent DNS name or IP address, since it is critical that this value be accessible to all RSS client readers. The default is the official name for the server and port number. It may be necessary to modify this value to account for different interfaces, DNS entries, or proxy servers.
- **Keep Report Copies.** Default 30. This is the number of copies of reports to keep on the system for each report. If the default value of 30 is used, then approximately one month of report data is retained when publishing daily, or approximately half a year of report data is retained when publishing weekly, or approximately 2 years of data is retained when publishing monthly.
- **Publish Report Versions.** Default "Newest". This value indicates how many RSS reports are listed in the channel. By default, only the last report published is listed as an RSS item, which means that, if the user does not access the report before it is again published, the user will not acquire the report via RSS. For example, if reports are published daily, and the user is unable to run the RSS reader for several days, that report data will be

unavailable to the user. Generally, the default value of "Newest" will be sufficient for most purposes.

- **Schedule Publishing.** Default "Daily". This value indicates how often reports are automatically published. If the setting is "Daily", then reports are published each night, at midnight. If the value is "Weekly", then reports are published at midnight Sunday morning. If the value is "Monthly", then the reports are published at midnight on the first of each month. In all cases, reports are published only after the newest reports have been created, so the reports contain the latest data available at the time they are published.
- **Use HTTP Authentication.** Default "No". This value adds HTTP authentication to the "rss" directory, so the user must login to the system with a valid username and password before any report can be downloaded. Most RSS aggregators will be able to handle this with no problems, although some aggregators may not support this function. By its nature, RSS is intended to be a way of distributing reports to the general public, which works against the idea of implementing HTTP authentication.

RSS Publishing Security Considerations

It is important to note that RSS is a way of distributing reports to the public, without requiring a login to CorreLog. This may have certain security implications, since it potentially exposes important network information to users.

By default the CorreLog configuration does not publish reports, and only a user with an "admin" type login can configure and enable publishing. To permanently disable RSS publishing, leave the default "Master RSS Enable" setting to "Disabled", and then remove or rename the RSS.exe program, residing in the sigma-web of the installation, to remove the tab from the screen.

You may also enable HTTP authentication, via the "RSS" screen, so that a user must have a CorreLog login to view any published report. You may also modify the "apache/conf/access.cnf" file of the system to limit the range of IP addresses that can access reports. You may configure an HTTP authenticating proxy, or firewall, so that only certain portions of your organization can view reports. Finally, you may selectively disable any report from being published, via the "Excel" screen, so that only certain reports are published, and others are kept private.

In particular, if you are publishing reports, it is probably a good idea never to include the actual messages in the report (as configured via the "Excel" screen) and generate only Summary data. Although this exposes IP addresses on the network, the actual message content is not published.

Support For ODBC Compliant Databases

In addition to generating Excel reports, CorreLog also supports updating ODBC relational databases with the same information. Users can update tables in MS-SQL, MS-Access, Oracle, or other databases by defining an ODBC "Data Source Name". This allows users to leverage existing report writers (including, but not limited to the powerful and inexpensive "Access" database) to generate reports on Message data, or run arbitrary SQL queries on this data.

To interface to an ODBC database rather than an Excel spreadsheet, the user creates a "System DSN" using the native Windows Administrative Tools > ODBC Data Sources dialog, normally located in the Windows Control Panel. This DSN is then configured with a username, password, and database name within the CorreLog "System > ODBC" tab.

If the user selects a System DSN rather than an Excel spreadsheet startpoint, then data is written to the specified database rather than the Excel spreadsheet.

In this case, the database is not published via RSS and the hyperlink for downloading the file is not operational. Instead, the user can access the database through an ODBC connection, using a third-party report-writing tool. The user can also use the CorreLog "System > ODBC" screen to view this data, or use any of the CorreLog Framework SQL functions documented in the companion "CorreLog Sigma Web Framework" Users Manual

Any tables with the named areas (listed above) are dropped and recreated in the named data source each time the report is generated. The specific table names are the same as the "named areas" of the Excel spreadsheet, documented in a previous section. This means that, if multiple reports are to be written to a database, each report must have its own database partition and ODBC DSN.

Configuring A DSN And ODBC Report Database

The specific steps needed to update a relational database rather than an Excel spreadsheet are documented as follows:

1. The user runs the ODBC Data Sources Dialog of the Windows platform, accessed via the Administrative Tools section of the Windows Control Panel.
2. The user clicks the "System DSN" tab of the dialog and clicks "Add" to add a new System DSN to start the Windows "Add New Data Source" wizard.
3. The user selects a driver for the database that will be updated. For example, the user can select Microsoft Access as the driver, which will be

available on all Windows platforms, including those that do not have Microsoft Office installed.

4. The user finishes the wizard, including any special configuration of the driver.
5. Within CorreLog, the user accesses the "System > ODBC" tab, and configures the "username", "password" and "usedb" fields for the system DSN created above.
6. Within the CorreLog "Reports > Excel" screen, the user creates a report, and selects as the "Startpoint" the name of the ODBC data source configured above. (This new data source will have been added to the dropdown list of startpoints.)

When the report generator runs (either at its regularly scheduled time, or manually) data will be loaded into the specified database rather than an Excel spreadsheet. The tables listed in the previous section will be created automatically and populated. Any existing tables with the same names will be dropped prior to creating the new tables.

Note that this technique permits multiple databases to be updated during a single run of the report generator. For example, the user might have several access databases, each with a separate DSN, and each filled with different message data and statistics. Similarly, there may be various partitions in a database that permit the user to join tables within SQL queries

Unique Table Prefix

If the report generator is loading data into a relational database, it may be useful to give the report a unique "Table Prefix" value, configurable on the "Report Edit" screen. This is a three-character value that is used as a prefix for any tables that the report generator creates, so that each report has uniquely named tables within the database.

This value will generally be required if multiple reports are loaded into a single database, because the CorreLog system will overwrite any existing table with a new table of information. If a table prefix is not specified, the database will contain information on only the last report that is generated (rather than information on each report generated.)

The Unique Table Prefix is never used by Excel spreadsheets, and is only if multiple reports are being loaded into the same database.

Database Table and Column Names

CorreLog automatically creates tables in the database, dropping any existing tables of the exact same name and recreating these tables each time the report generator runs.

It is important to note that the reporting facility does not actually preserve any data in the database, and that the database is simply a way of interfacing CorreLog to third party report writers. (It is therefore quite common to see the Microsoft "Access" database used as the report database, since it is found on all Microsoft systems, and is easily robust enough to support the fairly limited requirements of CorreLog.) The Administrator should be fully cognizant of the fact that CorreLog takes the liberty of dropping database tables before recreating them, hence permissions to drop tables in the database is a strict requirement of this system.

The specific table and column names are readily available for inspection using the "Reports > ODBC" tools, using standard SQL "select" statements. Each table has the same name as the "Named Area" of the Excel spreadsheet, listed in a previous section, with the addition of the three-character "Table Prefix" value (if any) that the user has configured in the Report Editor screen.

The SQL statements needed to select all the data from the various tables of the system are listed below.

Below, the value of (pfx) is the table prefix listed for the report, if any. For example, if the "Table Prefix" for a particular report is "all", then to select the report title, the statement that is entered into the ODBC tool is "select * from allReportTitle".

```
select * from (pfx)ReportTitle
select * from (pfx)ReportTime
select * from (pfx)DeviceSummary
select * from (pfx)UserSummary
select * from (pfx)FacilitySummary
select * from (pfx)SeveritySummary
select * from (pfx)FacilityData
select * from (pfx)SeverityData
select * from (pfx)KeywordData
select * from (pfx)MessageData
```

Each of the above select statements (given the appropriate value for "pfx") will retrieve all the data for a particular table.

Generally, most SQL databases are case sensitive; hence the table names (including the table prefix, if any) should be typed in the exact upper and lower case letters. The "select" and other SQL keywords are not case sensitive.

In addition to supporting the "select" command, the ODBC tool permits any other SQL query, useful for an "ad hoc" examination of report data. More information on database support can be found in the CorreLog "Sigma Web Framework" Users Manual.

Other Report Types

This section has concentrated on the Excel reporting facility, but it is worthwhile noting the other reporting capabilities of CorreLog.

More formal types of reporting are also available, in addition to Excel reporting. Each of these report types auto-execute at midnight, or on demand, and generate reports for users describing information subsequent to their time of execution. CorreLog supports the following additional reporting facilities.

- **E-Mail Reports.** The user can send raw-email messages (up to some arbitrary number, by default 500 messages) as an attachment to an e-mail message. This is a useful type of reporting when messages occur only occasionally, such as to support file integrity monitoring, logins to specific devices, or other anomalous types of messages. The user configures a particular thread, and optional match patterns that further qualify the message, precisely targeting specific message content.
- **Pivot Reports.** The user can analyze "regular" field limited data using the Pivot report capability. This type of report is especially useful for analyzing and reporting on web servers, firewalls, or other message sources where each message consists of regular columns of data. The user can define specific fields of different message types and the Pivot report generator will automatically parse these fields
- **Graph Reports.** The "Graphing" facility is provided in the "Reports" area, and can be used to capture message rates by minute, hour, and day. The user can configure specific "Graph" reports, save these reports, and then access the reports for printing. Most web browsers will provide reasonable printouts of graphs using the browser "Print" function. (Some browsers may require landscape mode to capture all the data.)

Finally, it is useful to note that each screen of CorreLog can be regarded as a separate printable report (available via the standard web browser "Print" function.) Because CorreLog will operate on an 800 pixel wide browser, all information on all types of screens can be captured on demand, saved, and / or

printed. This provides a useful way of generating report information on a variety of different screens for archiving, forensics, and analysis.

Pivot Reports

CorreLog Pivot reports, configured and accessed via the "Reports > Pivot" tab of the system, represent a particularly important class of reporting, the usage of which is beyond the scope of this document. The "Pivot Report Generation Function" is documented within a separate manual, which can be accessed from the "More > User Manuals" link of the CorreLog Server.

The "Pivot Report Generation Function User Guide" provides a detailed description of Pivot Reporting functions, their application and usage, as well as a comprehensive tutorial and procedures for creating these types of reports. For detailed notes on Pivot Reports, see the "More > User Manual" link, or contact CorreLog Support for the latest version of this manual in PDF format.

Section Summary And Additional Notes

1. CorreLog generates reports in Microsoft Excel format, and these reports can be downloaded via the CorreLog web interface.
2. Reports are normally generated each night at midnight. The user can also generate reports on demand via the "Run Report" button at the top of the screen, including individual generation of reports via the "Edit" screen of each report.
3. CorreLog permits generation of multiple reports, summarizing the message content, and acting upon all messages or messages in a particular CorreLog correlation thread.
4. Excel Reports permit the user to configure various options, and include the ability to parse messages and update the spreadsheet with message content.
5. To parse messages, the operator must first enable message reporting, which may have implications in data security (by exporting messages to a spreadsheet, which may then be distributed to other parties.) These considerations should be weighed when exporting data.
6. Including messages in an Excel spreadsheet may lengthen the time for the report to be generated, so that the report generation program can take many minutes.
7. In addition to downloading reports from the CorreLog web interface, an administrator can elect to publish these reports using RSS (Real Simple

Syndication) protocol, so these reports are publicly available to users on a daily, weekly, or monthly basis.

8. To customize the reports, a user can create a new Excel Startpoint Template, or modify the "Default.xls" file, to contain new graphs, formulas, macros, and formatting.
9. CorreLog supports ODBC compliant databases rather than Excel spreadsheets. The user creates a "System DSN" and configures a username, password, and database within the CorreLog "System > ODBC" screen.
10. When using a System DSN rather than an Excel spreadsheet as the report Startpoint, the data cannot be published or downloaded via the Web. Instead, the user can interface the database to a report writer or dashboard program.
11. When using a System DSN, the tables within the database are automatically created and populated. Any existing tables with the same name are dropped before any new tables are created. This requires each report to exist in its own database partition, and have a unique ODBC identifier configured for the report, or requires a unique Table Prefix for each report.
12. The "Pivot" reporting tool is used to parse regular data fields (such as that commonly found in firewall data, HTTP server logs, and many other messages.)
13. The "Pivot" report tool allows the user to quickly see the relationship between message counts for the field values within a message.
14. The "Pivot Report Generation Function" is documented within a separate manual, which can be accessed from the "More > User Manuals" link of the CorreLog Server

Appendix A: List of All Screens

This appendix provides a listing of screens of the CorreLog Server, and a brief description of each screen, useful as a site map.

The CorreLog Server contains all screens in the “CorreLog/sigma-web” directory, which forms a nested hierarchy of screens and functions. The operating theory behind this system is detailed and well documented in the CorreLog Web Framework User Manual, available under a different cover.

An administrator can rename and re-order this list of screens by creating new directories, and renaming the executable CGI programs that reside therein. An administrator and developer can easily add new screens to the system, ranging from simple HTML and text files to batch files to more sophisticated CGI programs.

Refer to the CorreLog Web Framework User Reference Manual, available in the “s-doc” directory of this site for notes on configuration, adaptation, and modifications that can be made to this system, including programming tools and reference material.

Home Screen

This is the first screen displayed on login to the CorreLog Server, and permits the user to download and read documentation and utilities. A version of this manual, in PDF form, is available for download from this page.

Dashboards Screen

This screen is displayed when the user clicks on the "Dashboards" tab. It allows the user to view and configure dashboards on the system. These dashboards consist of a layout file, and "gadgets" that depict real-time system status information.

Dashboard > Add New Dashboard Wizard

This screen is displayed with the user clicks on the "Add New" button at the bottom of a dashboard screen. The user can create a new dashboard for the system. (The user can also create a dashboard using the "Edit Layout" screen, described below, and using the "SaveNew" button on that screen.)

Dashboard > Edit Layout

This screen is displayed when the user clicks on the "Edit Layout" button at the bottom of a dashboard screen. The user can change the layout file, provide panel titles and hyperlinks to other dashboards, rename dashboards, save new dashboards, or edit gadgets for each dashboard panel.

Dashboard > Edit Gadget

This screen is displayed when the user clicks on the "Edit Gadget" button on the "Edit Layout" screen, and also is accessed via the "Edit Gadget" button in the upper right corner of each gadget windowpane. The user can specify a gadget and the gadget specific parameters for the dashboard panel.

Messages > Search

This is the first screen displayed when the user clicks on the "Messages" navigation tab. The screen provides a list of recent events, and permits the user to search the list for keywords, and page the list. This is the main entry point for the GenDex search engine.

Messages > Search > Advanced Search

This screen is accessed from the "Search" screen (above) by clicking on the "Advanced Search..." hyperlink at the top of the display. The screen permits the user to compose advanced searches based upon match patterns, exclude patterns, partial matches, severity and facility matches, as well as time of day.

Messages > Search > Device Information

This screen is accessed from the Search screen, as well as from various other screens in the system, when the user clicks on a device name hyperlink. The Device Information screen shows the IP address, DNS name, the ping response, and additionally shows SNMP values to help identify the IP address.

Messages > Catalogs > Devices

This screen shows the list of all devices that have send messages, the count of the messages since startup, and a count of the messages since the device was first detected. The screen provides a summary overview of the Devices catalog.

Messages > Catalogs > Devices > View Device Groups

This screen is accessed when the user clicks the "View Groups" link at the upper right of the "Devices" screen. The "View Device Groups" screen shows a list of all the address groups, and the rolled up status of each group, including counts, and time of last message for the group.

Messages > Catalogs > Devices > Device Item Catalog

This screen is accessed when the user clicks on one of the IP address hyperlinks in the above Devices Catalog screen, and shows the actual messages associated with a selected IP address. The user can delete the catalog item from this screen.

Messages > Catalogs > Devices > Device Catalog > Analyze

This screen is accessed when the user clicks on an IP address, to view the catalog of messages for the device, and then clicks the "Analyze" link in the upper right of the display. The screen permits the operator to analyze the messages and their occurrence counts for each Device, User, Facility, and Severity found in the list of messages.

Messages > Catalogs > Devices > Advanced Device Monitor Configuration

This screen is accessed when the user clicks on the "Advanced" button on the "Messages > Device" screen. The user can enter parameters that affect when the status icon colors turn red or green, and the severity of internal messages logged when devices are discovered, idled, or reactivated.

Messages > Catalogs > Users

This screen shows the list of all usernames discovered by CorreLog, and the count of messages for each user since startup, and the count of messages since the username was first detected. The screen provides a summary overview of the Users catalog.

Messages > Catalogs > Users > User Item Catalog

This screen is accessed when the user clicks on one of the username hyperlinks in the above Users Catalog screen, and shows the actual messages associated with a selected username. The user can delete the catalog item from this screen.

Messages > Catalogs > Users > User Catalog > Analyze

This screen is accessed when the user clicks on a user name hyperlink to view the catalog of messages for the username, and then clicks the "Analyze" link in the upper right of the display. The screen permits the operator to analyze the messages and their occurrence counts for each Device, User, Facility, and Severity found in the list of messages.

Messages > Catalogs > Users > User Discovery Configuration

This screen is accessed when the user clicks the "Advanced" button on the "Messages > Users" screen, and permits the user to configure match patterns, field positions, and exclude patterns needed to enable the discovery of new users on the system. These values are pre-configured for CorreLog, but this screen may be necessary to make adjustments or add new types of user discovery.

Messages > Catalogs > Facilities

This screen shows the list of all facilities that been specified in received messages, the count of the messages since startup, and a count of the messages since the facility was first detected. The screen provides a summary overview of the Facilities catalog.

Messages > Catalogs > Facilities > Facility Item Catalog

This screen is accessed when the user clicks on one of the Facility hyperlinks in the above Facilities Catalog screen, and shows the actual messages that are associated with the selected facility. The user can delete the catalog item from this screen.

Messages > Catalogs > Facilities > Facility Catalog > Analyze

This screen is accessed when the user clicks on a facility hyperlink to view the catalog of messages for the facility, and then clicks the "Analyze" link in the upper right of the display. The screen permits the operator to analyze the messages and their occurrence counts for each Device, User, Facility, and Severity found in the list of messages.

Messages > Catalogs > Severities

This screen shows the list of all severities that been specified in received messages, the count of the messages since startup, and a count of the messages since the severity was first detected. The screen provides a summary overview of the Severities catalog.

Messages > Catalogs > Severities > Severity Item Catalog

This screen is accessed when the user clicks on one of the Severity hyperlinks in the above Severities Catalog screen, and shows the actual messages that are associated with the selected severity. The user can delete the catalog item from this screen.

Messages > Catalogs > Severities > Severity Catalog > Analyze

This screen is accessed when the user clicks on a severity hyperlink to view the catalog of messages for the severity, and then clicks the "Analyze" link in the upper right of the display. The screen permits the operator to analyze the messages and their occurrence counts for each Device, User, Facility, and Severity found in the list of messages.

Messages > Aux

This screen shows the list of messages that have been filtered for today. The screen shows all the messages for all configured filters. The messages in all files are deleted each night at midnight and restarted. The user can search the list of filtered messages, and page through the list. Users can also report on these messages using the Excel and Reporting functions. (The "Parms" screen, discussed below, also controls one aspect of filtering, which is the "de-duplicate" filter time period.)

Messages > Config > Filters

This is the first screen displayed when the user clicks on the "Advanced" navigation tab, and depicts the list of filters that are applied to incoming messages. When a message is filtered, it is temporarily moved to the "Messages > Aux" screen (described above.) The user can add, modify, or delete filters. Filters can be based upon address, facility, device, keyword, and / or time of day.

Messages > Config > Overrides > Address

This is the first screen displayed when the user clicks on the "Overrides" navigation tab. The screen shows a list of "Address Overrides". The user can override the name of the device in a message, before the message is logged, useful for handling network address translation of IP addresses, and useful in other cataloging applications.

Messages > Config > Overrides > Address > Parse Specification

This screen is accessed when the operator clicks on the "Address Overrides" navigation tab, clicks "Advanced" and then clicks the "Auto-Override With Parse Specification" button. The screen allows an administrator to configure and test a parse specification that will automatically override the device in the message with some portion of the message content, useful when using a syslog collector, load balancer, network address translation, and other strategies.

Messages > Config > Overrides > Facility

This screen shows a list of "Facility Overrides". The user can override the facility of a message, before the message is logged, using this screen. A user specified facility replaces the actual facility, based upon the device name, the facility, the severity, a keyword, and / or the time of day. This is useful for cataloging messages, and executing action programs.

Messages > Config > Overrides > Facility > User Facility Editor

This screen is accessed when the operator clicks on the "User Defined Facility List" hyperlink on the above "Facility Override" screen. This happens to be the most deeply nested screen within the suite of CorreLog screens. The User Facility Editor permits an operator to define a new facility that can then be selected in one of the other screens, such as the "Correlation Threads" or "Actions" screen, to supply a higher level of organization to event messages.

Messages > Config > Overrides > Severity

This screen shows a list of "Severity Overrides". The user can override the severity of a message, before the message is logged, using this screen. A user specified severity replaces the actual facility, based upon the device name, the facility, the severity, a keyword, or the time of day. This is useful for cataloging messages, and executing action programs.

Messages > Config > Overrides > Text

This screen shows a list of "Text Overrides". The user can override a particular field or value within a message by blanking out the value. This permits sensitive data such as passwords, credit card numbers, other security items from being logged or displayed.

Messages > Config > Overrides > Text > Advanced Configuration

This screen is accessed by clicking on the "Advanced" button of the "Text Override Editor" screen, and permits the user to specify several advanced override functions, such as user name masking, and lists of specific keywords that are automatically masked by the system.

Messages > Config > Forwarding

This screen allows the operator to configure forwarding of messages or any auxiliary file to a third-party syslog collector, useful for multi-tier management strategies, or backing up the data for failover operation. Although other ways exist to forward data, this screen provides the fastest way to relay data to another syslog collector or CorreLog server.

Messages > Config > Colors

This screen allows the user to associate or change the colors of messages that appear in various locations in the CorreLog suite of screens, based upon the logged message severity.

Messages > Config > Parm

This screen allows the user to configure specific parameters that affect the performance of CorreLog processes, such as the location where Syslog data is stored, the location of the Net-SNMP software, and the de-duplication filter time value.

Correlation > Threads

This screen shows the list of all correlation threads that been configured by the user. These are special, arbitrary catalog items that permit the user to obtain lists of messages based upon the time of day, device, facility, severity, and keyword matches. The user can add, modify, and delete match patterns and specifications from this screen.

Correlation > Threads > View Thread Groups

This screen is accessed when the user clicks the "View Groups" link at the upper right of the "Threads" screen. The "View Thread Groups" screen shows a list of all the thread groups, and the rolled up status of each group, including counts, and time of last message for the group.

Correlation > Threads > Correlation Item Catalog

This screen is accessed when the user clicks on one of the Correlation Thread hyperlinks in the above Correlation screen, and shows the actual Syslog messages that are associated with the user defined thread. The user can delete the catalog item from this screen, or from the previous screen.

Correlation > Threads > Thread Catalog > Analyze

This screen is accessed when the user clicks on a "Thread" hyperlink to view the catalog of messages for the thread, and then clicks the "Analyze" link in the upper right of the display. The screen permits the operator to analyze the messages and their occurrence counts for each Device, User, Facility, and Severity found in the list of messages.

Correlation > Triggers

This screen permits the user to set system flags when specific messages are received. These flags can then be used to qualify the correlation "Threads" and "Actions". This provides a fundamental way of correlating specific information, and creating contexts for messages

Correlation > Actions

This screen shows a list of user configurable actions that occur when certain messages are logged. The user can configure program names and arguments that are executed when messages match a particular address, facility, severity, message keyword, and / or time of day.

Correlation > Config > Address Groups

This screen allows the user to configure device groups that can be used in the "Match IP Address" fields of the correlation "Threads" and "Actions" screens. The user can configure lists of IP addresses and wildcards that represent groups of devices.

Correlation > Config > Thread Groups

This screen allows the user to configure thread groups that can be used in to organize the "Threads" screen. The user can configure multiple thread groups, each group consisting of a match expression that matches the title of zero or more threads. The thread groups can be viewed via the "Correlation > Threads" screen..

Correlation > Config > Macros

This screen allows the user to configure macro expressions that can be used in the "Match Expression" fields of the correlation "Threads", "Triggers", and "Actions" screens. The user can define a macro that represents a complex match pattern, and then use this macro (with other macros) to match messages.

Correlation > Config > Lists

This screen allows the user to configure macro expressions that can be used in the "Match Expression" fields of the correlation "Threads", "Triggers", and "Actions" screens. These are similar to the "Macros" (discussed above) but instead consist of long lists of items, any which of can be used to match or exclude a message. This allows the operator to configure whitelists or blacklists.

Correlation > Config > Templates

This screen is a wizard that permits the user to load, merge, and replace existing correlation configurations and rules, or save the existing rules to a file. This allows users to checkpoint existing configurations, or quickly switch the system to a new configuration.

Alerts > Counters

This screen permits the user to configure thresholds on any system counter, so that message rates can be detected. The user can select from various system counters. The Alert screen feeds information back into the Syslog server process, so that it can be further correlated. The Alert screen can also be used to open tickets on the system.

Alerts > Devices

This screen is similar to the "Alerts Counters" screen above, but implements CorreLog "instance management", so that messages will create new instances of alerts, useful for situations where alerts are to be tracked on a "per device" basis. (See section on "Active Instances" within this manual, for more information.)

Alerts > Patterns

This screen provides general utility in detecting when certain patterns of messages have been received, based upon the state of "Triggers" on the system. The screen operates in a fashion similar to the "Alerts" screen, sending Syslog messages when patterns of messages are detected, and assigning tickets to system users.

Alerts > Custom

The "Custom Alert" screen extends the range of the alerting facility to include execution of arbitrary alerting programs. These external programs are launched at schedule intervals. The output of the alerting program is read by CorreLog, compared to a user define match expression, and a threshold applied to the number of matches can open a ticket.

Alerts > Config > Alert Formulas

This screen provides advanced functions that work with the "Alerts > Counters" screen, permitting the user to create formulas that reference multiple system counters, and run math expressions to calculate a single result, which can then be alarmed with a threshold.

Alerts > Config > Auto-Learn

This screen controls the various parameters of the CorreLog "auto-learn" function, which automatically adjusts alert threshold each night based on received message history, thereby minimizing the number of internal CorreLog alerts and open tickets.

Tickets > Opened

This screen shows is the first screen displayed when the user clicks on the "Tickets" tab. The screen shows the currently opened tickets on the system. Tickets are opened by the "Alerts" component, and represent the highest level of correlation on the system.

Tickets > Opened > View Ticket Groups

This screen is accessed when the user clicks the "View Groups" link at the upper right of the "Tickets > Opened" and "Tickets > Closed" screens. The "View Ticket Groups" screen shows a list of all the ticket groups, and the rolled up status of each group, including counts, and time of last ticket for the group.

Tickets > Closed

This screen shows the closed tickets on the system. The user can close individual tickets, or close all tickets associated with a user or ticket group. The closed tickets are retained on the system for the "Keep Data" interval configured by the user, by default 30 days.

Tickets > Actions

This screen shows a list of user configurable actions that occur when certain tickets are opened, closed, or modified. The user can configure program names and arguments that are executed based upon ticket information (similar to the Correlation > Actions screen described above, but reflecting the state of tickets.)

Tickets > Config > Ticket Groups

This screen permits the user to create ticket assignees, which represent groups of tickets. The user configures the alert facility to open tickets and assign these tickets to CorreLog users, or to ticket groups. This organizes the tickets into groups.

Tickets > Config > Parm

This screen provides miscellaneous parameters that affect ticket execution, including a master enable, throttles to limit ticket action rates, and the ability to enable automatic de-duplication of tickets opened on the same day.

Reports > Query

This screen permits the user to run a complex query against raw data, returning a list of match results that can be further graphed, searched, and analyzed. The screen performs similar functions to the "Messages > Search" screen, except can perform more complex (and time consuming) searches, especially useful for forensics. The screen searches Log Data, Archive Data, Auxiliary Files, Tickets, and External data.

Reports > Query > Run

This screen is accessed via the "Query" hyperlink in the upper right of each CorreLog screen, or when the user clicks "Run Report" on the "Reports > Query" screen. The operator can specify the complex patterns associated with a particular query across a range of messages.

Reports > Query > Saved Queries

This screen is accessed via the "Saved Queries" hyperlink, and contains a list of queries that have been saved (on the above "Query Run" screen, by specifying a name for the query.) The operator selects the saved query, and can modify or delete the query.

Reports > Query > Analyze

This screen is accessed when the user clicks on the "Analyze" hyperlink, and permits the operator to analyze the query results and their occurrence counts for each Device, User, Facility, and Severity found in the list of messages. The Query Analyze screen also provides a special "Parse" function to allow the execution of a parse expression (to see certain fields within the query results.)

Reports > Audit > User Activity

This screen allows the operator to access audit reports on general user activity, reporting on all users contained in the "Messages > Users" tab of the system. An operator can configure one or more report viewers to view this data, and view the reports in HTML, CSV, and text format.

Reports > Audit > Device Activity

This screen allows the operator to access audit reports on general device activity, reporting on all managed devices contained in the "Messages > Devices" tab of the system. An operator can configure one or more report viewers to view data, and view the reports in HTML, CSV, and text format.

Reports > Audit > Perimeter

This screen allows the operator to access audit reports on perimeter activity, reporting on raw messages that contain two or more IP addresses, where one of the IP addresses is external to the organization. An operator can configure one or more report viewers to view this data, and view the reports in HTML, CSV, and text format.

Reports > Audit > Account Management

This screen allows the operator to access audit reports on account management activity, reporting on messages related to account creation, deletion, modification, group assignments, and other changes to LDAP and Active Directory dealing with user access and authentication. An operator can configure one or more report viewers to view this data, and view the reports in HTML, CSV, and text format.

Reports > Audit > Tickets

This screen allows the operator to access audit reports on system tickets generated by CorreLog alerts. This provides supervisory visibility to the ticketing system and all incidents identified by the CorreLog Server program. An operator can configure one or more report viewers to view this data, and view the reports in HTML, CSV, and text format.

Reports > Audit > Score Cards

This screen allows the operator to configure and view "Score Cards", which are reports that reflect the number of messages received given a specific requirement. Score Cards are especially useful for demonstrating compliance to regulatory requirements and standards, but also are useful for self assessment, as to whether the system is configured for correct coverage of organizational and operational units.

Reports > Excel

This screen allows users to create or modify Excel reports, and download reports directly from the CorreLog web interface after login. Reports are created each night. The user can generate reports on demand from this screen.

Reports > E-Mail

This screen allows users to create or modify E-mail reports, which consist of summary information and attachments of e-mail message lists. Reports are created each night. The user can generate reports on demand from this screen.

Reports > RSS

This screen allows users to configure the RSS publishing interface, which permits RSS aggregators to collect report information without a login to CorreLog, scheduled on a daily, weekly, or monthly basis.

Reports > ODBC

This screen is a simple interface to the Windows ODBC facility. The user can configure the username and password values associated with an ODBC "Data Source Name" (DSN), and also run simple SQL statements. (The screen is mentioned here, but not discussed in any detail within this current manual. See the "CorreLog Sigma Web Framework" Users Manual more information on this screen.)

Reports > Graphs

This screen permits the user to graph message counts, and depict either daily or hourly event rates. The user can specify complex match patterns to limit the message counts to a range of messages.

Reports > Pivot

This screen permits the user to analyze "regular" data, such as firewall data, web server data, or other data that consists of regularly occurring fields. These reports parse the data and create graphs of field items, useful for determining things such as most active URLs, most active source destinations, most active firewall rules, etc.

System > Prefs

This screen is the first screen displayed when the user clicks the "System" navigation tab, and shows the user preferences of the system. The user can configure various parameters (such as initial login screen) that apply to the individual CorreLog user. This screen may also be accessed via several links within CorreLog, including the username link in the far upper right of the display.

System > Logins > Users

The screen allows the administrator to configure logins, passwords, and profiles for the various CorreLog users. Various options apply. These are the local logins to the server, documented in previous sections, and apply to both HTTP and Web Screen type logins

System > Logins > Profiles

This screen allows the administrator to configure access profiles that can then be selected by the administrator when creating a user login. Profiles are used to limit the visibility and access of the system to certain tabs, dashboards, and pinned items.

System > Logins > Security

This screen allows the administrator to configure various extra security options for the system, including SSPI (Active Directory) authentication, whether the user is to login to the system using HTTP or Web Screen login dialogs, and various items related to local password expiration and lockout.

System > SMTP

This screen permits the user to configure SMTP server parameters that are used by the actions facility (i.e. the SENDMAIL program) as well as the E-Mail Report facility. The administrator configures standard SMTP values such as server IP address, and authentication type.

System > ParmS

This screen allows the user to adjust certain System level parameters, such as the CGI timeout, and tab color values. This screen is documented in detail within the "CorreLog Sigma Web Framework" Users Manual.. (The screen is mention here, but not discussed in any detail with this current manual.)

System > Schedule

This screen allows the user to adjust system startup parameters, and allows the user to specify other programs to be run on a scheduled basis. The screen modifies the CO-svc.exe configuration file. This screen is documented in detail within the "CorreLog Sigma Web Framework" Users Manual. (The screen is mention here, but not discussed in any detail with this current manual.)

More > User Prefs

This screen is accessed via the "More" drop down menu at the top right of the display. The screen is identical to the "System > Prefs " screen, permitting the user to modify his or her preferences.

More > User Links

This screen is accessed via the "More" drop down menu at the top right of the display. The screen permits the user to augment the list of links in the program header with other links of interest. These links appear only in the user's private login session.

More > Keywd Index

This screen is accessed via the "More" drop down menu at the top right of the display. The screen displays the list of message keywords, identical to the list displayed via the "Keyword Index" link on the "Messages > Search" screen.

More > Send Msg

This screen is accessed via the "More" drop down menu at the top right of the display. The screen permits the user to send a Syslog message to CorreLog, identical to the "Post Message" link on the "Messages > Search" screen.

More > Lookup SID Tool

This screen is accessed via the "More" drop down menu at the top right of the display. The screen provides general utility in translating a Microsoft Security Identifier (SID) into readable text. (SID values may occur in messages, and are typically represented in a format such as "S-1-5-N-N", which translates into a user name, account name, or some other human readable object.

More > Expr Tool

This screen is accessed via the "More" drop down menu at the top right of the display. The screen allows a user to enter complex match expressions (including global variables and macros) and see the results of the match expression compared to a user supplied text string. This screen is documented in the "CorreLog Advanced Correlation System User Guide", included as a standard manual in every CorreLog package. . (The screen is mention here, but not discussed in any detail with this current manual.)

More > Geo IP Tool

This screen is accessed via the "More" drop down menu at the top right of the display. The screen provides general utility in determining the Geo-location of an IP address. The Geo IP database is used by a variety of CorreLog screens, including the "Reports > Audit > Perimeter" reports.

More > Site Map

This screen is accessed via the "More" drop down menu at the top right of the display. The screen displays a list of all screens (similar to this list) and links to access any screen in the system.

More > User Manuals

This screen is accessed via the "More" drop down menu at the top right of the display. The screen displays a list of all user documentation that comes with the system. This includes the documents on the "Home" screen, as well as other manuals on the system

More > Extensions

This screen is accessed via the "More" drop down menu at the top right of the display. The screen displays a web folder of extra software components and extensions that support the CorreLog system, in particular McAfee EPO extension software, but also site-specific information.

More > Support

This screen is accessed via the "More" drop down menu at the top right of the display. The screen launches the CorreLog Support screen in a separate window, directly connecting to the CorreLog public website.

More > Sys Info

This screen is accessed via the "More" drop down menu at the top right of the display. The screen displays a status of the system, including program version, processor info, disk space info, and other values. (This screen is always available, even in the license to execute expires.)

Appendix B: Syslog Protocol

Over the long and often frenetic history of networked computing technology, many methods for the management of computer networks have come and gone. Methods such as CMIP and TN1 are no longer considered viable. Some methods, such as WBEM have never entered mainstream management techniques. But some of these protocols have tenaciously remained, and Syslog Protocol is one of them. Today, it is one of the most ubiquitous, simplest, most interoperable, and least appreciated of the management techniques available to administrators.

Syslog is a formal specification for sending messages from one computer system to another. It is formally documented in RFC 3164, published in 2001. The protocol itself has been in use since 1980. The invention of the protocol is generally attributed to Eric Allman, inventor of the Unix sendmail program, (as well as the founder of the most popular convention for indenting C-language programs, known as the “Allman Style” of bracketing.)

Unlike many management protocols (including SNMP) the format of the message is quite simple. The message is sent in clear text to a UDP port, normally port 514. The message contains two essential fields.

1. **Numeric Header.** The Syslog message includes a simple text prefix, encoded in ASCII, consisting of an integer number bracketed by “<” and “>” characters. The integer number is between zero and 191.
2. **Text Message.** This message header is followed by a text string of less than 1024 characters in length. This field can be further broken down into fields, although this is not essential or consistently enforced.

Within RFC 3164, an attempt is made to further codify the specification, by defining the format of embedded time strings, identifying sub strings and conventions with which to compose the text message portion.

Although this is worth attention, none of this is practical to consider. The above two rules comprise virtually the ONLY aspects of the message specification that can be completely relied upon for all Syslog message generation. In fact, this single paragraph describes all that is needed for a programmer to get started in generating Syslog messages.

RFC 3164: The BSD Syslog Protocol

It is of interest to review RFC 3164. It begins with the following sentences: “Since the beginning, life has relied upon the transmission of messages. The messages may signal danger, the presence of food, or the other necessities of life, and many other things.”

The RFC continues on to later note: “One of the fundamental tenants of the Syslog protocol is its simplicity. This simplicity has greatly aided the acceptance and deployment of Syslog.”

To a large degree, when these two sentences are extrapolated, they tell the entire story of Syslog protocol. It was intended to permit one-way transmission of messages, and has survived because of its extremely simple nature.

Given that fact, RFC 3164 can be discounted heavily in any design considerations regarding Syslog senders and receivers. The evidence is that the protocol is so trivial to implement that an RFC to describe that protocol in depth is almost overkill. To send a Syslog message, a sender opens a connection to a machine, where a process is listening at UDP port 514. The user sends a clear text message in the format described previously (i.e. a numeric header and some text), and then disconnects and continues without requiring acknowledgement.

The above paragraph is not meant to discourage a reader from reviewing the Syslog RFC. In addition to being a beautiful summary of the protocol, it is also one of the more engaging and humorous RFCs that an engineer is likely to encounter. For example, at one point, the RFC diverges from a technical discussion of security considerations and describes how messages, in the form of sex pheromones, can lure innocent animals into death traps. The RFC cites “Chemical Mimicry” as one if its thirteen different references. At another point, the RFC discusses cooking:

“An odor may be considered to be a message that does not require any acknowledgement. People tend to avoid bad odors but are drawn to odors that they associate with good food. The acknowledgement of the receipt of the odor

or scent is not required and indeed it may be the height of discretion to totally ignore some odors. On the other hand, it is usually considered good civility to acknowledge the prowess of the cook merely from the ambiance wafting from the kitchen.”

Forwarding Syslog Messages

One of the most important aspects of Syslog protocol, which has given it such leverage and longevity, is that Syslog messages are designed to be easily forwarded to a centralized Syslog collector. This is in addition to collecting Syslog messages on a local platform. The Syslog collector can be located on a main server in the enterprise (such as the CorreLog program) to aggregate all the messages of a particular facility source and / or priority.

The ability to forward Syslog messages is at the foundation of the protocol. All meaningful systems that implement Syslog support it. In fact, the facility and severity codes, so ubiquitous with the CorreLog Server, were originally intended to permit filtering and forwarding of messages based both on their interest level and pertinence, but also based upon which administrative station would be interested in aggregating certain messages, such as mail administrators, or UUCP administrators.

It may be that a message is forwarded several different times before it reaches its destination. Obviously, information at one level of the enterprise may have no pertinence, while the same information (such as in the hands of a network administrator or application specialist) might be extremely useful. RFC 3164 discusses this as follows:

“Flexibility was designed into this process so the operations staff would have the ability to configure the destination of messages sent from the processes running on a device. In one dimension, the events that were received by the Syslog process could be logged to different files and also displayed on the console of the device. In another dimension, the Syslog process could be configured to forward the messages across the network to the Syslog process on another machine.

“The Syslog process had to be built network-aware for some modicum of scalability, since it was known that the operators of multiple systems would not have the time to access each system to review the messages logged there. The Syslog process running on the remote devices could therefore be configured to either add the message to a file, or to subsequently forward it to another machine.”

Without this ability, there could be no system like the CorreLog program. Administrators would have to either gather information from each machine, or build software to relay information themselves, such as with remote shell or RCP

calls. However, given that messages are easily forwarded, it makes sense to forward them, a fact so trivial and obvious that it seems peculiar that this is not commonplace in all enterprises!

What makes this data aggregation slightly difficult is that the volume of Syslog messages can be tremendous, so that the scalability of the central Syslog server is quickly brought to its limit. Previously, the ability to search through these messages, catalog these messages in significant ways, has become too cumbersome for average software systems. This is precisely the role that the CorreLog Server is intended for, and the basic problem that the system answers. The CorreLog Server is meant to be the single data aggregation point for huge volumes of Syslog messages.

Syslog Message Content

One of the difficulties of Syslog protocol, which distinguishes it from SNMP and other request / response protocols, is that the end user has no control over the message content. If there was some certification procedure that was required in order to participate in the Syslog development process, this might not be too aggravating. However, any designer (either clever or inept) can generate Syslog messages that eventually end up in the CorreLog Server. RFC 3164 elaborates on this issue as follows:

“The writers of the operating systems, processes and applications have had total control over the circumstances that would generate any message. In some cases, messages are generated to give status. These can be either at a certain period of time, or at some other interval such as the invocation or exit of a program. In other cases, the messages may be generated due to a set of conditions being met. In those cases, either a status message or a message containing an alarm of some type may be generated. It was considered that the writers of the operating systems, processes and applications would quantify their messages into one of several broad categories.”

In practice, it appears that most Syslog messages are very pertinent. The CorreLog software can handle those messages that are either erroneous or badly formed as follows:

1. If the message facility code or the severity of the message seems inappropriate, based upon the message content or the various other factors of the message, the user can configure the system to override the facility or severity code, thereby making the message more proper. The end user has complete control over these two fields.
2. If no particular Facility seems to match the message, for example the message is in a special class that is not covered by the RFC 3164 facility codes, the user can actually define a new facility that makes sense in the

context of the enterprise, and then override the normal facility to be the new facility.

3. If the message content is not fixable or patchable, or has no pertinence whatsoever, and is just wasting bandwidth, the user can elect to simply filter the message out at very input of the system, with no other processing taking place on the message. The filtered messages are still collected and logged (for a time) but are purged from the system at a higher rate. In the case of the CorreLog Server, filtered messages are retained ONLY for 24 hours maximum.

The above features of the CorreLog Server greatly enhance the practicality of using Syslog protocol by increasing the end user control over the message.

Syslog Message Header: Facilities and Priorities

Throughout this manual, references have been made to the Facility and Severity codes, with the good assumption that the reader understands what these terms mean. However, given the importance of these two arguments of all Syslog messages, it seems worthwhile to provide a formal definition. This definition is found in RFC 3164, and is elaborated upon here.

Message facilities and severities are numerically coded with integer decimal values. Each facility and severity, in addition to being given a fixed number, is also associated with a formal “name”.

The facility names originate from system level Unix processes, and were intended to identify the actual name of the process generating the Syslog event. Later, this evolved to include identification of related functionality (such as “mail”, or “kernel”). Many of the facility codes are deprecated, and good candidates for being “overridden” as discussed in the previous section. In particular the “uucp” facility has been almost entirely deprecated. (Herein it is replaced with the name “network”). The “ftp” and “ntp” facilities are well on their way to being deprecated, although are still in some use within organizations, especially those with Unix based operations.

In a similar fashion, severity values were selected to indicate how messages were to be prioritized on the system. These severity values are sometimes referred to a “priorities” (although the use of this term may be confusing, since the “priority” is also sometimes referring to the facility / severity combination.) Unlike facility numbers, which have no ordinal value, severities have a definite rank, ranging from 0, the most important and pertinent severity, to 7, which is the lowest and most easily discarded message.

Confusion regarding Severities and Priority names abound. This document summarizes the actual severity and priority values used by the CorreLog Server, along with a detailed explanation. Refer to the following two tables.

Facility Codes And Their Meaning

The basic facilities, defined by RFC 3164, are discussed below.

Kernel	0	These are messages related to the Unix kernel process, or generated by very low-level driver software and system programs.
User	1	These are typically user-defined messages. This facility is used (and over-used) as a central way of defining messages that have not been otherwise classified.
Mail	2	These are messages related to the SMTP system, Microsoft Exchange, as well as mail relay systems, and sometimes e-mail programs.
System	3	This is another catchall type of facility that is often over-used, but generally related to system services, Unix daemons not otherwise classified. It can also indicate
Security	4	These are messages related to security processing, such as login detection, virus protection, and intrusion detection systems. Other security related messages are found in the "Audit(13)" and "Alert(14)" facilities.
Internal	5	Originally, these were messages related only to the internal operations of the Syslog protocol process, but have evolved to include general internal processes, often related to performance monitoring, but occasionally simply the internal workings of the system.
Printer	6	These messages are related to the Unix "lpd" process, and can also indicate problems with printer hardware, printer queues, and other types of queues that are not particularly related to printers.
News	7	These are messages related to the Network News processes. This facility is sometimes used to indicate low severity news events, such as a system being brought down.
Network	8	These messages are related to network events. This facility corresponds to the Unix "uucpd" process, which has been fairly well deprecated in the age of high-speed TCP networks. (UUCP is an acronym for "Unix to Unix Copy".)

Lock	9	This facility is listed in the RFC as “clock”, but is often renamed as “lock”, and used for locking mechanisms, such as file locking queues. It is often substituted for the “Clock(15)” facility code. It may (on some systems) be identical to the Clock(15) facility.
Auth	19	This is similar to the “Security(4)” facility, but is generally reserved for authorization errors, such as invalid logins. It is somewhat synonymous with both “Security(4)” and “Audit(13)”. It represents one of the areas of the RFC that is not clearly delineated, hence is subject to interpretation.
FTP	11	These messages are related to the Unix “ftpd” process, and FTP program, which is somewhat deprecated but still in use. This facility is sometimes used for non-FTP protocol messages related to file transfers.
NTP	12	These messages are related to the Unix “ntpd” (News Transport Protocol) processes. This is somewhat deprecated, but can still be found on a variety of Unix platforms.
Audit	13	This is similar to the “Security(4)” and “Auth(19)” facility codes, but mainly appropriate for audit processing, including performance monitoring. For example, a performance monitor might use this facility to periodically send the disk space and disk utilization statistics to the Syslog process for data collection. The messages that use this facility should be pertinent to performance reporting.
Alert	14	This is a general-purpose (hence heavily overused) facility to indicate an Alert condition. This may be somewhat confusing, because this is really a severity rather than a facility. Ideally, these messages would represent problems with the alerting process rather than actual alerts.
Clock	15	These messages are related to the Unix clock daemons, and other processes involved with time synchronization and maintenance. This facility is also sometimes used to mark event times, such as by issuing a Syslog message via the Unix “cron” or Windows “at” program. Some scheduler programs use this facility. Occasionally, due to ambiguities in the RFC, this facility is confused with, and substituted for the “Lock(9)” facility.

Local0	16	This is a user definable facility, used by Cisco and many other vendors. It is of the used in application software, and is an ideal candidate for being modified by the CorreLog Server to provide a more meaningful facility name, based upon the message content.
Local1	17	This is another user definable facility. See notes regarding the Local0(16) facility.
Local2	18	This is another user definable facility. See notes regarding the Local0(16) facility.
Local3	19	This is another user definable facility. See notes regarding the Local0(16) facility.
Local4	20	This is another user definable facility. In particular, this is commonly used by RedHat clustering software, and is used by the Cisco PIX software, and is used in some Perl scripts. See notes regarding the Local0(16) facility for more information.
Local5	21	This is another user definable facility. See notes regarding the Local0(16) facility.
Local6	22	This is another user definable facility. See notes regarding the Local0(16) facility.
Local7	23	This is another user definable facility. See notes regarding the Local0(16) facility.

For users wishing to define their own Syslog messages, in particular Syslog messages generated by application programs, the Local0 through Local7 facilities are good candidates in the absence of other obvious alternatives.

Some firewalls and routers permit a user to associate messages with specific facilities, in which case the Local0 through Local7 facilities may be dedicated to a router or firewall configuration, based upon some arbitrary enterprise policy. The CorreLog program can provide a more meaningful and descriptive facility name through a user defined facility that overrides one (or all) of the Local0 thru Local7 standard facilities.

If a user elects to use the Local type facilities, these messages should have unique content, which makes it easy to filter and override. In practice, many developers will attempt to wedge their message facility into one of the other more commonly defined facilities, perhaps in the mistaken belief that their message pertinence will be increased. Experience demonstrates that it is better to define these facilities in the enterprise through some sort of written policy or operational guidelines, thereby promoting consistency (at least within the organization.)

Severity Codes And Their Meaning

The basic severities, defined by RFC 3164, are discussed below.

Debug	7	This represents the lowest severity, and is reserved strictly for debugging the system. In practice, debug messages can be totally ignored by everyone. It is common to over-use this severity, and users should always consider promoting this severity to “Info” in a production environment.
Info	6	These are informational messages, which can be reviewed later (having some pertinence) but which can be operationally ignored because they have no effect on management activities.
Notice	5	These are messages that are sent with the intention of being noticed. They have a fairly significant level of importance. It is common to filter these messages (wrongly) since they are meant to be noticed and not discarded. Hence, a filter should generally not remove any messages with this severity.
Warning	4	This is a significant message. It should be taken seriously and reacted to, because it signifies a non-trivial degree of risk. There may not be any corrective action needed with this type of message.
Error	3	This is a highly significant message. The message indicates that corrective action, manual intervention, or operational change is necessary. All error severities will require acknowledgement and a fix or system configuration change. No error message should ever be ignored.
Critical	2	A critical situation exists that requires immediate attention. All other activities should be set aside and the problem addressed as soon as possible. Possible risk to security or data or infrastructure is eminent.
Alert	1	An extremely critical condition exists that will require immediate intervention by the highest levels of management, requiring whatever resources necessary to immediately fix. Data has been lost, security has been breached, or infrastructure has been damaged.
Emergency	0	This severity should NEVER be used. Human safety is in jeopardy, immediate evacuation is needed, police and / or firefighters need to be called.

Given the obvious nature of the common meaning for these severities, it is startling to see how many times the severities are actually misinterpreted or ignored, both by operations management, and the designers of Syslog messages.

For example, because “Emergency” is the highest severity, many naïve designers may decide to arbitrarily use this severity for messages that are rightly “Info” type messages, knowing that an “Emergency” message will probably not ever be filtered. This can be extremely annoying. If this situation occurs, it is easily corrected in the CorreLog Server, by using the “Severity Override” feature of the system to adjust the severity to its proper place.

From The Viewpoint Of Information Theory

One way to consider the usage of facility and severities is through an informational perspective, as follows: The more often a user receives a message, the less pertinent it is, and the less information that is actually imparted by the message. For example, if a program sends a message often, at periodic intervals, the most significant thing that can occur is for the program to STOP sending messages. In this case, the program is using bandwidth ineffectively. It would be better to filter those messages at their source.

On the other hand, if a message occurs very infrequently, that message is highly significant, regardless of the message content. The very fact that a particular message occurred, and could not be anticipated, means it has achieved a high degree of information content. This property of information is called “entropy”, which is actually the unit of measure for information systems. The more entropy that exists in an information stream, the more random the information, and the more pertinent the information actually is. As entropy increases, so does information content.

The problem occurs when entropy (i.e. the amount of randomness in an information system) achieves the same level of entropy as the environment in which the information is processed. In that case, no further information can be perceived. The information is too random. This point – at which information can no longer be useful – is called the “bandwidth” of the information system.

Knowledge of this fact can permit a high degree of information tuning, through the facility overrides, match patterns, and filters. The CorreLog Server should be configured, wherever possible, to accept what is unique and reject what is common. Where a message can be made MORE unique (such as by applying an override) then that message becomes more interesting to everyone and more valuable to the enterprise.

Security Concerns

The final parts of RFC 3164 are concerned heavily with the security of Syslog protocol. Various security considerations are discussed, including potential hacks and attacks. Syslog is based upon UDP protocol, which is agreeable to denial of service attacks. UDP is a protocol that is not reliable (i.e. it is possible that Syslog messages are dropped when buffers become full.) Additionally, Syslog messages are not encrypted hence can be sniffed off the network fairly easily.

Both of these security concerns can be addressed fairly easily with relay programs that use TCP protocol and encryption. For example, the CorreLog program provides such a mechanism (if it is needed) in the form of the CO-bridge.exe program, available on request. This is a simple UDP tunnel that encrypts data, and uses a request / response mechanism to insure reasonable delivery of messages. Other utilities exist that may also be useful, such as SSL tunnels commonly available from the web.

In most situations, because Syslog messages will be transmitted on secure networks, security considerations are not overbearing. As is commonly mentioned, if an organization is under attack, the Syslog messages will probably be the least concern to the attacker. (This rule may be modified depending upon the degree of importance that the Syslog protocol has assumed in the organization, and the awareness of this fact by the attacker.)

Conclusions

Syslog protocol is a highly valuable, well tried, and typically under utilized method of achieving stable networks, informed managers, and problem awareness.

Because it relies on the copious existing Syslog generators on the network (in the form of Unix boxes, firewalls, application programs, routers) and because Syslog is easily adapted to homegrown monitoring, it is a highly cost effective way of implementing a management solution or augmenting a management strategy.

RFC 3164 concludes with the following cautionary statement, the truth of which should be easily recognized and accepted:

“There have been attempts in the past to standardize the format of Syslog messages. Many good thoughts have come from that effort, and interested implementers may want to find notes or papers regarding that effort. It must be cautioned that the simplicity of the existing system has been of tremendous value to its acceptance. Anything that lessens that simplicity may diminish that value.”

Appendix C: Dashboard Gadgets

The "Dashboards" screen is the entry point to the CorreLog dashboard facility, which permits the user to display real-time data about various elements of the system, such as message rates, top devices, top users, and many other data items.

Each dashboard consists of various window panels. Each window panel can contain a user specified "Gadget". The particular gadgets are specified and configured via the "Edit Layout" button at the bottom of the screen, or by clicking the "Edit Gadget" icon at the upper right of each window panel.

The user can create, modify, and delete dashboard configurations. The operator can make this the default login screen, and select a default dashboard, using the "User Preferences" screen of the system.

Specific gadgets are as follows:

Analyze-Thread-Gadget

The Analyze-Thread-Gadget allows the operator to see the "Devices", "Users", Facilities", and / or "Severities" associated with a selected thread during a user selected time interval. The gadget operates in a fashion similar to the "Analyze" screen available on Catalog screens. This gadget may be better suited to a scrolling type window, but can work in any dashboard pane.

Custom-Accounting-Gadget

The Custom-Accounting-Gadget allows the user to select certain counters (for example, "All Messages", "All Tickets", or thread counters) and provide a label for these counters, useful for accounting and audit purposes. The gadget also permits the user to specify their own custom style sheets, header file, and footer file. This gadget is often used in "multi-tenant" mode, so that a tenant of the system can account for the messages pertaining to their particular area of interest.

Custom-Alerting-Gadget

The Custom-Alerting-Gadget allows the user to select certain alerts (via match expressions applied to the alert text message. The gadget then shows the status of these alerts rolled up to a single item. The gadget requires a match pattern that matches the alert text for "Counter", "Pattern", "Device", "Custom" and other alerts. Once configured, user can drill down into gadget to see the various matched alerts on the system and their status. Additionally, the gadget permits the user to custom style sheets, header file, and footer files, useful for advanced customization of the gadget.

Device-Rate-Gadget

The Device Rate Gadget allows the operator to see a list of recent devices, along with the rate of change (either positive or negative) in the message rate for the device during a span interval of time ranging from 1 to 15 minutes. The gadget is useful for seeing what devices are no longer sending messages or have experienced a recent surge or drop in message rates.

External-URL Gadget

The External-URL Gadget allows the operator to specify an arbitrary web page segment to be displayed in a dashboard panel. The operator specifies the URL, including a required "http://" prefix, as a gadget parameter via the gadget editor screen. When the dashboard is displayed, the panel will show the contents of the specified URL, clipped at the corners of the panel. This provides a method of easily extending the dashboard facility to include external information, pictures, forms, and other components at an external URL.

Forwarding-Gadget

The Forwarding Gadget allows the operator to inspect the state of a single CorreLog forwarder by "Message Source" ("Msg01", Aux01", etc.) as a method of determining the current state of the forwarding operation. The gadget depicts the forwarding parameters and current counts for successful forwards, forwarding errors, and messages bypassed for the forwarder due to a non-match of the configured match facility and severity. The operator can configure this dashboard gadget for each forwarder of interest. The actual Forwarding operation is configured on the "Message > Config > Forwarding screen.

Gauge-Alert-Gadget

The Gauge-Alert-Gadget monitors a correlation alert on the system. Any alert defined in the "Alerts > Counters" facility can be specified. The gadget creates a gauge indicating the alert threshold, alert state and current value of the alert. This permits the operator to easily inspect the state of any alert on the system, and whether the alert is near to exceeding, or has exceeded, its configured threshold.

Gauge-Loading-Gadget

The Gauge-Loading-Gadget can be used to monitor the system loading to assess performance and assist with system tuning. The gadget monitors several different types of performance, selected by the operator, as follows: the "Daily Message" loading displays the amount of daily messages received, compared to the maximum licensed limit; the "Device Process Delay" loading monitors the delay between message reception and the update of the "Devices" screen; the "Correlation Process Delay" loading monitors the delay between message reception and the update of the "Threads" screen; the "Queued Actions" load monitors the number of actions pending execution.

Gauge-Ticket-Gadget

The Gauge-Ticket-Gadget monitors a "ticket user" on the system, as defined in the "Tickets" facility. Any ticket user can be specified as an argument to the gadget, including the "All" user. The operator specifies the ticket user name, a threshold, and a time interval. This permits the operator to easily inspect the number of open tickets assigned to various users on the system, as an indication of current system activity.

Graph-Message-Rate-Gadget

The Graph-Message-Rate-Gadget displays the rate of received messages on the system during the last five minutes, the last sixty minutes, the last five hours, or the last ten days. The operator can select either a vertical bar graph, or a line chart (where the line chart permits higher resolution, but requires Flash or Java.) This gadget allows the operator to easily inspect short-term or long-term message rates on the system as a method of determining message activity. Note that, depending upon the general loading of the system, some display modes and graph types may take slightly longer to render. If this is a consideration the user may wish to try different time intervals for the graph (for example a five minute graph type instead of a 60 minute graph type.)

Graph-Thread-Rate-Gadget

The Graph-Thread-Rate-Gadget displays the rate of received messages for a particular correlation thread on the system during the last five minutes, the last sixty minutes, the last five hours, or the last ten days. This gadget is similar to the "Graph-Message-Rate-Gadget" described above, but allows the operator to monitor the activity of a single thread (as opposed to all messages.) The operator must select a valid thread as part of the gadget configuration. Any thread defined in the "Correlation > Threads" facility can be specified. Experience shows this type of gadget is fairly ubiquitous and used quite often in dashboard depictions.

Graph-Thread-Delta-Gadget

The Graph-Thread-Delta-Gadget displays the rate of received messages for two different threads, and depicts the difference between these two graphs. This is useful if two threads are generally synchronized with each other, and the difference between two different threads is significant. The operator must select one or two valid threads as part of the gadget configuration. Any thread defined in the "Correlation > Threads" facility can be specified. This graph is quite specialized; most users will want to use the "Graph-Thread-Rate" gadget instead of this one. (See notes above.)

Last-Message-Gadget

The Last-Message-Gadget allows the operator to specify up to eight different correlation threads. For each thread, the gadget depicts the last message logged by the thread since midnight. (If no message exists for the thread, no value is displayed.) This gadget permits the operator to view recently received messages, especially useful for viewing message that occur only occasionally (such as messages that reflect anomalies and exceptions.) This type of gadget is especially useful for "Scroll" or "Map" type layout files.

Map-Ticket-Gadget

The Map-Ticket-Gadget allows the user to create a visual representation of processes, devices, or applications by specifying an image file and then locating "ticket groups" on the image. The operator specifies a URL to an image file, and then specifies the name and XY coordinates (on the map file) of up to eight different ticket groups. If any ticket is opened for one of the ticket groups (during the current day), the icon color for the group appears "red", otherwise the icon color appears "green". Options to assist in the design and configuration of the gadget appear at the bottom of the "Edit" screen.

Parse-Thread-Gadget

The Parse-Thread-Gadget is similar to the Analyze-Thread-Gadget, but allows the operator to specify a parse function to display the unique counts of a certain field parsed from each message within a thread (as opposed to the device names, user names, etc., contained within the thread.). The operator selects an existing Thread identifier, an optional match, and the parse specification. The dashboard then parses the field from the message thread and tabulates counts of each unique field. Parse specifications are the same as used in other CorreLog locations, such as the "Query" and "Excel" report functions.. This gadget may be used in any windowpane, but is particularly well suited for a "Scroll" type layout file.

Severities-Gadget

The Severities-Gadget displays a breakdown of the severities of all messages on the system since Correlog program startup. (Note that this list of severities reflects data only since CorreLog was first started, and not necessarily the severity of all messages on the system.) The operator can specify the severity list as a pie chart, a stack chart, or a table of fields. This permits the operator to easily inspect the severity of recent messages received on the system.

Text-HTML-Gadget

The Text-HTML-Gadget provides general utility in allowing an operator or administrator to configure a window panel with arbitrary text or HTML, such as to add annotations, user supplied images, or hyperlinks to the display. The operator can enter arbitrary HTML and text to a maximum of 400 characters. If the text does not fit in the window panel, it is automatically clipped at the window pane edges.

Top-Alerts-Gadget

The Top-Alerts-Gadget depicts the state of the top five alerts on the system, indicating whether the alert is above or below its threshold, as well as the current alert value, and time interval for the alert. First, any alerts that are currently set (red) are listed, followed by alerts that are clear (green). To limit the list of alerts to a specific range, the operator can specify a match keyword or expression as part of the gadget configuration, such as "Facilities", "Devices", or some logical combination of keywords. This provides a method of limiting the display to a range of system alerts, or five specific alerts.

Top-Devices-Gadget

The Top-Devices-Gadget depicts the state of the top five devices on the system. These top devices can be the most recently updated devices, or the devices that have received the most messages since CorreLog system startup. The operator can limit the list of devices to a specific "Device Group" (configured in the "Correlation > Config > Address Groups" screen.) This provides a method of limiting the display to a range of devices, or five specific devices on the system. The operator can select the display to be a horizontal bar chart, a pie chart, a stack chart, or a table of fields.

Top-Facilities-Gadget

The Top-Facilities-Gadget depicts the state of the top five syslog facilities on the system. These top facilities can be the most recently used facilities, or the facilities used most often since CorreLog system startup. Facilities may be the standard syslog facilities, or user defined facilities (configured in the "Messages > Config > Overrides > Facilities" screen. The operator can depict facilities as a horizontal bar chart, a pie chart, a stack chart, or a table of fields.

Top-ODBC-Gadget

The Top-ODBC-Gadget can be used to query values directly from an ODBC capable database. The user composes an arbitrary SQL query that returns at least two columns, where the first column contains labels, and the second column contains values. If the second column is numeric, the user can graphically depict the top five results as a horizontal bar chart, a pie chart, or a stack chart. If the second column is not numeric, the "fields" display mode can depict arbitrary data. Note that the user specified SQL query will typically include an "order by" clause to return "top" values.

Top-Threads-Gadget

The Top-Devices-Gadget depicts the state of the top five correlation threads on the system. These top threads can be the most recently updated threads, or the threads that have received the most messages since CorreLog system startup. The operator can limit the list of threads to a specific set of threads using a keyword or expression qualifier that matches the thread title. This provides a method of limiting the display to a range of threads, or five specific threads on the system. The operator can select the display to be a horizontal bar chart, a pie chart, a stack chart, or a table of fields.

Top-Ticket-Users-Gadget

The Top-Ticket-Users-Gadget depicts the state of the top ticket users on the system, based upon the number of open tickets for each user, the number of closed tickets, or all tickets. This provides an easy way to view ticket activity on the system in terms of who is being assigned (or has been assigned) tickets. The operator can limit the list of ticket users to a specific set using a keyword or expression qualifier. This provides a method of limiting the display to a range of ticket users, or five specific users on the system. Note that the state of individual ticket users can also be monitored using the "Gauge-Ticket-Gadget" described previously.

Top-Users-Gadget

The Top-Users-Gadget displays the top user references on the system. These top users can be the users that have most recently sent messages, or the users that have generated the most messages since CorreLog system startup. This provides an easy way to view user activity on the system in terms of who is generating messages. The operator can select the display to be a horizontal bar chart, a pie chart, a stack chart, or a table of fields.

Watch-Device-Gadget

The Watch-Device-Gadget displays the message status of a single system device, showing a depiction similar to the "Messages > Devices" screen. The operator specifies the "Idle Time Threshold". If no messages have been received from the specified device after that time, the background fields of the gadget turn red, furnishing an easy way to view message activity for a device. The operator can configure multiple "Watch Devices" on a single dashboard to watch activity from critical system devices.

Additional Notes on Gadgets

1. Gadgets are stand-alone programs that reside in the "s-cgi/gadgets" folder. This folder is usually protected via an "htaccess.txt" file (requiring logins to access gadget information.)
2. Gadgets can be externally accessed (depending on web server permissions) by a static URL available via the "View External URL" link available on the editor for a gadget. This permits an administrator to display a gadget on a web page outside of CorreLog.
3. Most gadgets have additional options that allow the administrator to suppress drill down into the gadget, and allow the administrator to customize the error message (if any) associated with the gadget. This is available via the "Additional Options" link toward the bottom of the editor for a gadget.
4. Dashboard layout files can be customized, and new dashboard layouts created by modifying the contents of the "dash/layouts" directory of the system.
5. Dashboard configurations can be programmatically created or externally modified by modifying the contents of the "dash/config" directory of the system.
6. Installation defaults, useful for restoring dashboard configurations, are contained in the "dash/install" directory of the system.
7. The latest help on gadget (which may be different herein, depending upon whether a service pack has been installed at the CorreLog Server) is generally available via the "Gadget Help" link, accessed on the dashboard editor screen. The online help will usually reflect the latest gadgets and their functions, especially for intermediate versions of the system.

CorreLog is pleased to accept and discuss enhancement requests that pertain to new functionality, or any existing functionality of its dashboard system. The dashboard facility is intended to be driven by customer requirements, and is often updated to reflect the current needs and ergonomic considerations of the CorreLog user database. Feel free to contact CorreLog support to begin a discussion.

Appendix D: Windows Process List

This section contains a list of the standard and optional windows processes that may appear in the task manager of the CorreLog Server system. All persistent CorreLog processes begin with a "CO-" prefix, making it easier to identify CorreLog processes. The information here includes the standard CorreLog programs, as well as optional CorreLog adapters.

CO-Action.exe

This process runs the Ticket, Correlation Actions, and Custom Alerts of the program. The process is a standard CorreLog program that should appear in the Task manager of all CorreLog Server systems.

CO-Apache-Tls.exe

This process is the secure CorreLog Apache TLS program, installed as part of the CorreLog Enhanced Encryption package. The process replaces (or augments) the standard CO-Apache.exe program that comes with the default version of CorreLog. Two of these processes should be executing on the system.

CO-Apache.exe

This process is the default standard Apache server that comes with the default version of CorreLog. Two of these processes should be executing on the system, unless the user has disabled this feature (such as to use the CO-Apache-TLS.exe process, above.)

CO-Assoc.exe

This process runs the "Association Monitor" of CorreLog. The process is a standard part of the GSIP version of CorreLog, or added to the default version of CorreLog via the co-x-x-x-assoc.exe package. This process should continuously execute for the "Correlation > Associations" screen to function properly.

CO-Catlog.exe

This process runs the Correlation "Threads", "Alerts", and "Tickets" functions of the CorreLog server. The process is a standard CorreLog program that should appear in the Task manager of all CorreLog Server systems.

CO-Devlog.exe

This process runs the Message "Devices", "Facilities", and "Severities" screens of the CorreLog server. The process is a standard CorreLog program that should appear in the Task manager of all CorreLog Server systems.

CO-Fmon.exe

This process runs the CorreLog File Integrity Monitor Agent, an optional agent program that can be downloaded and installed on Windows (or UNIX) platforms to continuously check file integrity. The program can be downloaded and installed from the "Home" screen of the standard CorreLog Server.

CO-Gendex.exe

This process runs the Generate Index function, which continuously indexes incoming messages and maintains the "logs/dex" folders of the CorreLog Site. This program is required to support the CorreLog high-speed search function, and is standard CorreLog program that should appear in the Task manager of all CorreLog Server systems.

CO-Ipv6.exe

This process is the CorreLog "IPv6 Gateway" process. The software is optionally installed on CorreLog platforms to support IPv6 protocol.

CO-Maint.exe

This process is launched each night at midnight, and is responsible for performing periodic maintenance of the system, generating reports, creating archive files, and limiting the number of files on the system. The program is a transient process, but may take several hours to execute on busy systems.

CO-Ping.exe

This process runs the CorreLog Ping Adapter, and is added to a standard version of CorreLog via the "co-X-X-X-ping.exe" package. The process continuously pings network devices, and sends notifications when devices fail to respond.

CO-Queue.exe

This process runs the CorreLog File Transfer Queue Adapter, and is added to a standard version of CorreLog via the "co-X-X-X-queue.exe" package. The process waits for files to be written to a directory, and then sends the files to CorreLog.

CO-Sess.exe

This process runs the "Session Monitor" of CorreLog. The process is a standard part of the GSIP version of CorreLog, or added to the default version of CorreLog via the co-x-x-x-sess.exe package. This process should continuously execute for the "Correlation > Sessions" screen to function properly.

CO-Snmp.exe

This process runs the CorreLog SNMP Adapter, and is added to a standard version of CorreLog via the "co-X-X-X-snmp.exe" package. The process continuously polls network devices with SNMP "get" requests, and compares values to thresholds. When a threshold is violated, the process sends a notification to CorreLog.

CO-Sqlmon.exe

This process runs the SQL Table Monitor Adapter, and is added to a standard version of CorreLog via the "co-X-X-X-sqlm.exe" package. The process polls an SQL table for new entries (based on an SQL query), and then sends notifications to CorreLog containing the new table entry.

CO-Svc.exe

This process is the main service manager for the "CorreLog Framework" service. The process launches other programs, and also runs the "Schedule" screen of CorreLog. The process is a standard CorreLog program that should appear in the Task manager of all CorreLog Server systems.

CO-Syslog.exe

This process listens for Syslog messages at the standard UDP port of 514, and logs these messages to the "logs" directory. The process is responsible for filtering and overrides. The process is a standard CorreLog program that should appear in the Task manager of all CorreLog Server systems. If this process is not running, then no data is logged, including data logged by most adapters and plug-in components.

CO-Sysmsg.exe

This process is the Windows Agent program, which monitors the event logs and streaming log files of the system. The process is a standard CorreLog program that should appear in the Task manager of all CorreLog Server systems. Additionally, The program can be downloaded and installed on other platforms from the "Home" screen of the standard CorreLog Server.

CO-Systrap.exe

This process listens for SNMP Trap messages at the standard UDP port of 162, and logs these messages to the "logs" directory. The process is added to a standard version of CorreLog via the "co-X-X-X-trap.exe" package.

CO-Trecv.exe

This process is the tunnel receiver program, which is a standard part of the CorreLog server, but not generally enabled except at those sites using the Windows Tools Set "Tunneling" features, as documented in the Windows Tool Set Manual.

CO-Wmi.exe

This process runs the CorreLog WMI Adapter, and is added to a standard version of CorreLog via the "co-X-X-X-wmi.exe" package. The process continuously polls network devices with WMI requests to acquire log data, which is logged to the "logs" directory.

For Additional Help...

Detailed specifications regarding the CorreLog Server, add-on components, and resources are available from our corporate website. Test software may be downloaded for immediate evaluation. Additionally, CorreLog is pleased to support proof-of-concepts, and provide technology proposals and demonstrations on request.

CorreLog, Inc., a privately held corporation, has produced software and framework components used successfully by hundreds of government and private operations worldwide. We deliver security information and event management (SIEM) software, combined with deep correlation functions, and advanced security solutions. CorreLog markets its solutions directly and through partners.

We are committed to advancing and redefining the state-of-art of system management, using open and standards-based protocols and methods. Visit our website today for more information.



CorreLog, Inc.

<http://www.CorreLog.com>

<mailto:support@CorreLog.com>

Alphabetical Index

A

Access / 107 132 154 155 156
Account / 72 140 146 173
Action, Adding Correlation At Web Interface / 105
Action, Creating New Program / 103
Action, Detailed Program Usage And Debugging / 107
Action, Detailed Steps In Configuring User / 100
Action, Testing / 106
Action, Using Existing Program / 101
Actions, Correlation Environmental Variables / 103
Actions, Ticket Facility / 124
Active / 96 140 170 173
Activity / 139 172
Adapter / 19 201 202
Adapters / 18
Adapters, Installing CorreLog / 18
Adding / 105 123
Adding, Manually Tickets / 123
Adding Correlation Action At Web Interface / 105
Additional Correlation References / 96
Additional Forwarding Techniques / 73
Additional Notes / 21
Address Groups / 92
Addresses / 52 53 70
Adjusting / 71

Admin / 29
Administrative / 14 17 75 109 154
Administrator / 14 72 109 133 156
Administrators / 115 132 138 181
Advanced Search Term Notation / 49
Advisories / 36 37
Advisories, Auditing Recent / 37
Advisories, Local System / 36
Advisory / 36 37 38
Agent / 17 18 70 200 202
Aggregation / 9
Aggregators / 137
Alert, Related Messages And Source Definition / 123
Alert Auto-Learn Function / 94
Alerting / 82 93
Alerting Functions / 93
Alerts, Defining Correlation And Tickets / 30
Allman / 179
Alone / 6
Alphabetical Index / 205
Analyze / 11 63 64 129 138 163 164 165 168 172 191
Analyze-thread-gadget / 191 195
Apache / 40 199
Applications / 9 39
Applications, CorreLog / 39
Apply / 51 59 60
Archive / 36 47 76 129 171
Archive, CHKSC Program – Checking Security Codes / 76
Archived / 47 77
Archiving / 6 9 39 75 76 77
Archiving, Data Functions / 75
Archiving, Data Message Digest Functions / 76
Archiving Filtered Data / 77
Area / 156
Areas / 148
Arguments / 106
Ascii / 43 146 179
Asfexe / 126
Asn1 / 114
Assign / 121
Assigned / 26 121 122
Assignee / 121 129 132
Assignee, Ticket / 121
Assistance / 103
Associated / 53 55 123
Association / 200

- Associations / 200
- Audit / 37 72 77 123 137 138 139 150 151 172 173 176 185
- Audit, Change Capability / 77
- Audit Report Types / 139
- Audit Reports / 138
- Auditing / 28 37
- Auditing Recent Advisories / 37
- Auth / 185
- Authentication / 153
- Auto- / 10 53 133 166
- Auto-Close Duplicate Tickets Setting / 127
- Auto-Learn, Alert Function / 94
- Auto-Masking User Names For Data Privacy / 71
- Auto-Overriding Device Names With Parsed Values / 70
- Auto-close / 127 128
- Auto-learn / 6 10 30 94 95 127 133 135 170
- Auto-learning / 133
- Auto-mask / 71 72
- Auto-masking / 71 72
- Auto-override / 70 71
- Auto-overriding / 70
- Auto-ranging / 60
- Automated / 39
- Automatic / 9
- Automatically / 116
- Automatically Opening Tickets / 116
- Automation / 99
- Automation, CorreLog / 99
- Aux-01 / 73
- Aux-1 / 35 73
- Aux-16 / 73
- Aux-8 / 35
- Aux01 / 193
- Aux Message Files / 67
- Auxiliary / 34 35 67 73 171

B

- BSD Syslog Protocol / 180
- Background / 9 108
- Basic / 14 23 82 120 150
- Basic, CorreLog Usage / 23
- Basic CorreLog System Installation Steps / 14
- Bios / 102 117
- Browsing / 39
- Button / 95

C

- C-language / 179
- CHKSC Program – Checking Archive Security Codes / 76
- CMDB Table Definition / 130
- Capability / 10 77
- Capability, Change Audit / 77
- Cards / 140 173
- Case / 87
- Catalog / 46 61 62 63 64 65 78 94 163 164 165 168 191
- Catalog, Deleting Item / 64
- Catalog Information / 46
- Catalog Message And History Counters / 62
- Cataloging / 9 62
- Catalogs / 9 61
- Catalogs, Data / 61
- Category / 152
- Caution / 57
- Center / 10
- Change / 77
- Change Audit Capability / 77
- Changed / 124 125
- Changing Syslog UDP Port Configuration / 43
- Checking / 76
- Checking, CHKSC Program – Archive Security Codes / 76
- Checkout / 16
- Chemical / 180
- Chksc / 76
- Cisco / 7 11 17 186
- Class-b / 92
- Clear / 51 87 97
- Clearing / 51
- Clearing Search Screen / 51
- Click / 20 28 29 31 32 34 35
- Clicking / 29 94 96
- Clinking / 63
- Clock / 185
- Close / 123
- Closed / 26 29 122 123 124 125 128 171
- Closing / 123
- Closing And Deleting Tickets / 123
- Cmdb / 129 130 131
- Cmdbexe / 130
- Cmip / 179
- Co-actionexe / 40 108 199

Co-apache-tlsexex / 199
 Co-apacheexex / 199
 Co-assocexex / 200
 Co-catalogexex / 100 107
 Co-catlogexex / 40 46 64 67 78 100 104 107 108 112 200
 Co-cmdbexex / 129
 Co-devlogexex / 40 200
 Co-fmonexex / 200
 Co-gendexexex / 40 200
 Co-ipv6exex / 200
 Co-maintexex / 41 47 75 200
 Co-pingexex / 201
 Co-queueexex / 201
 Co-sessexex / 201
 Co-snmpexex / 201
 Co-sqlmonexex / 201
 Co-svcexex / 40 41 48 175 201
 Co-syslogexex / 16 40 41 43 44 45 47 51 67 78 202
 Co-sysmsgexex / 7 11 202
 Co-systrapexex / 44 51 67 78 202
 Co-trecvcnf / 45
 Co-trecvexex / 19 20 43 44 45 103 202
 Co-trecvlog / 45
 Co-tsendcnf / 45
 Co-tsendexex / 20 44 45
 Co-uninstexex / 21
 Co-wmiexex / 202
 Code / 76
 Codes / 68 69 76 184 187
 Codes, CHKSC Program – Checking Archive Security / 76
 Codes, Creating User Defined Facility / 69
 Codes, Defining New Facility / 68
 Codes, Deleting User Defined Facility / 69
 Codes, Facility And Their Meaning / 184
 Collection / 6 39
 Color / 75
 Colors / 167
 Column / 156
 Combinations / 84
 Command / 133
 Commit / 142
 Common / 129
 Common Management Database Setup / 129
 Compliance / 24
 Compliant / 154
 Compliant, Support For ODBC Databases / 154

- Component / 82
- Components / 82
- Computers / 16
- Computers, Configuring Syslog On Network / 16
- Comspec / 104
- Concepts / 120
- Concerns / 189
- Conclusions / 189 189
- Configuration / 10 31 43 45 56 66 85 86 142 151 164 167
- Configuration, Changing Syslog UDP Port / 43
- Configuration, Dashboard And Layout Files / 56
- Configuration, Excel Report Items And Settings / 142
- Configuration, Other Syslog Files / 45
- Configuration, RSS Items And Settings / 151
- Configuration Of Correlation Rules / 85
- Configure / 42 43 54 68
- Configuring / 16 17 31 32 33 100 142 154
- Configuring, Detailed Steps In User Action / 100
- Configuring DSN And ODBC Report Database / 154
- Configuring E-Mail Reports / 33
- Configuring SMTP Notifications / 31
- Configuring Syslog On Network Computers / 16
- Configuring Ticket E-Mail And Notifications / 32
- Confirm / 138
- Connect / 84
- Connection / 83 84
- Connectivity / 14
- Considerations / 153
- Considerations, RSS Publishing Security / 153
- Content / 42 182
- Content, Syslog Message / 182
- Conventions / 41
- Conventions, Syslog File Naming / 41
- Copies / 152
- CorreLog, Basic System Installation Steps / 14
- CorreLog, Installing Adapters / 18
- CorreLog, Installing WTS On Windows Platforms / 17
- CorreLog, Referencing Gadgets Externally / 58
- CorreLog, Uninstalling Entire Framework / 20
- CorreLog Applications / 39
- CorreLog Automation / 99
- CorreLog Basic Usage / 23
- CorreLog Correlation / 81
- CorreLog Fast Start / 10
- CorreLog Installation / 13
- CorreLog Interactive Usage / 47

- CorreLog Reporting / 137
- CorreLog Server Features / 8
- CorreLog Tickets / 119
- Correlated / 74
- Correlation, Adding Action At Web Interface / 105
- Correlation, Additional References / 96
- Correlation, Configuration Of Rules / 85
- Correlation, CorreLog / 81
- Correlation, Defining Alerts And Tickets / 30
- Correlation, Defining New Threads / 26
- Correlation, Introduction To / 8
- Correlation, Reducing And Message Load / 34
- Correlation Actions Environmental Variables / 103
- Correlation Rule Dependencies / 96
- Correlation Rules And Rule Systems / 83
- Correlations / 77
- Correlog- / 110
- Correlog-database / 110
- Count / 28 35 130
- Counter / 30 93 192
- Counters / 62 82 96 97 131 169 170 193
- Counters, Catalog Message And History / 62
- Counting / 144
- Creating / 29 69 103 122
- Creating New Action Program / 103
- Creating User Defined Facility Codes / 69
- Critical / 93 106 139 187
- Current / 130
- Currentcount / 130 131
- Currentseverity / 130 131
- Custom / 148 170 192 199
- Custom-accounting-gadget / 192
- Custom-alerting-gadget / 192

D

- DN Helper Program / 53
- Daily / 54 59 60 153 193
- Dashboard / 56 58 135 162 191 198
- Dashboard Configuration And Layout Files / 56
- Dashboard Gadgets / 191
- Dashboard Screens / 56
- Dashboards / 56 162 191
- Dashboards Screen / 162
- Data, Archiving Filtered / 77
- Data, Auto-Masking User Names For Privacy / 71

- Data Archiving Functions / 75
- Data Archiving Message Digest Functions / 76
- Data Catalogs / 61
- Database / 129 130 139 154 156
- Database, Common Management Setup / 129
- Database, Configuring DSN And ODBC Report / 154
- Databases / 109 154
- Databases, Support For ODBC Compliant / 154
- Databases, Updating SQL ODBC / 109
- Date / 42 59
- Daycount / 130 131
- Days / 143 144 150
- Dayseverity / 130 131
- Db admin / 122
- De-duplicate / 67
- De-duplication / 67
- Debug / 33 127 187
- Debugging / 107
- Debugging, Detailed Action Program Usage And / 107
- Default / 152 153
- Defaultxls / 143 148 159
- Define / 9 148
- Defined / 68 69 166
- Defined, Creating User Facility Codes / 69
- Defined, Deleting User Facility Codes / 69
- Defining / 26 30 68
- Defining Correlation Alerts And Tickets / 30
- Defining New Correlation Threads / 26
- Defining New Facility Codes / 68
- Definition / 26 123 130
- Definition, CMDB Table / 130
- Definition, Related Messages And Source Alert / 123
- Delay / 193
- Delete / 69 142
- Deleting / 64 69 123
- Deleting, Closing And Tickets / 123
- Deleting Catalog Item / 64
- Deleting User Defined Facility Codes / 69
- Dependants / 85
- Dependencies / 96
- Dependencies, Correlation Rule / 96
- Dependents / 96 98
- Description / 54 130 131 152
- Detailed / 18 47 100 107 138 139 203
- Detailed Action Program Usage And Debugging / 107
- Detailed Steps In Configuring User Action / 100

Detection / 82
Device-rate-gadget / 192
Device, Auto-Overriding Names With Parsed Values / 70
Device, Managed Identification / 52
Device, SNMP Information / 54
Device Information Screen / 53
Devicesummary / 148
Devmsgs / 9
Dhcp / 52 53 68 70
Dialog / 154
Different / 24
Digest / 76
Digest, Data Archiving Message Functions / 76
Directive / 20
Directory / 43 107 140 173
Directory, Relocating Syslog Log / 43
Disable / 94
Disabled / 152 153
Discovery / 164
Disk / 14
Display / 130 131
Distributed / 10
Distribution / 151
Distribution, RSS Report System / 151
Dnhelperexe / 54
Dnshelperexe / 54
Domain / 53 54
Download / 18
Duplicate / 127 128
Duplicate, Auto-Close Tickets Setting / 127

E

E-Mail, Configuring Reports / 33
E-Mail, Configuring Ticket And Notifications / 32
E-Mail, Sending Messages / 112
E-mail / 32 33 34 38 112 150 157 173 175
Editing / 122
Editing Tickets / 122
Editor / 28 31 69 75 120 156 166 167
Emergency / 106 187 188
Enable / 71 72 73 94 129 152 153
Enabling / 43
Enabling Syslog TCP Reception / 43
Encrypted / 103
Encryption / 199

Encryptkey / 45
Engine / 48
Engine, GenDex Search / 48
Enhanced / 199
Entire / 20
Entire, Uninstalling CorreLog Framework / 20
Environmental / 103
Environmental, Correlation Actions Variables / 103
Eq-notice / 51
Equal / 52
Eric / 179
Error / 93 106 187
Errors / 85
Errorseverity / 44
Even / 151
Event / 7
Events / 33 139 141
Examples / 50 89 90 106
Examples, Search Term Notation / 50
Excel, Further Notes On Spreadsheets / 149
Excel Report Configuration Items And Settings / 142
Excel Report Message Parsing Rules / 145
Excel Reports / 140
Except / 44 87
Exchange / 69 184
Exclude / 113
Exclusion / 72
Executable / 133
Execute / 15 111 115
Existing / 101
Existing, Using Action Program / 101
Experienced / 41
Explorer / 18
Expr / 28 176
Expression / 87 89 90 91 98 143 169
Expressions / 87 88 91
Expressions, List Macro / 91
Expressions, Match / 87
Extensions / 177
External / 76 132 171 198
External-dash / 58 59
External-dashcnf / 58
External-url / 192
External, Tickets URL / 132
Externally / 58
Externally, Referencing CorreLog Gadgets / 58

Extra / 53 55
Extract / 18
Extraction / 48

F

Facilcnf / 46
Facility, Creating User Defined Codes / 69
Facility, Defining New Codes / 68
Facility, Deleting User Defined Codes / 69
Facility, Ticket Actions / 124
Facility Codes And Their Meaning / 184
Facilitydata / 149
Facilitysummary / 149
Failures / 29
False / 45
Fast / 10
Fast, CorreLog Start / 10
Feature / 59
Feature, Graph Summary / 59
Features / 8 20
Features, CorreLog Server / 8
Feed / 20 152
Feeddemon / 137 151
Feel / 198
Field / 88 145
File / 7 33 35 41 45 46 47 57 127 200 201
File, Log Report Information / 47
File, Syslog Naming Conventions / 41
Files / 45 56 67 171
Files, Aux Message / 67
Files, Dashboard Configuration And Layout / 56
Files, Other Syslog Configuration / 45
Filtcnf / 45
Filter / 34 35 36 66 67
Filtered / 9 67 77 79
Filtered, Archiving Data / 77
Filtering / 9
Filters / 35 38 66 73 79 166
Filters, Incoming Message / 66
Finish / 27 31
Firewall / 35 92 107
Flash / 194
Flexibility / 181
Forensics / 24
Formulas / 170

- Forward / 73
- Forwarding / 72 73 74 79 167 181 193
- Forwarding-gadget / 193
- Forwarding, Additional Techniques / 73
- Forwarding, Message / 72
- Forwarding Syslog Messages / 181
- Framework, Uninstalling Entire CorreLog / 20
- Frequently / 143
- From Viewpoint Of Information Theory / 188
- Front / 6
- Function / 63 94 133 138 158 159
- Function, Alert Auto-Learn / 94
- Functions / 75 76 77 93 145
- Functions, Alerting / 93
- Functions, Data Archiving / 75
- Functions, Data Archiving Message Digest / 76
- Functions, Special Message Parsing / 145
- Further Notes On Excel Spreadsheets / 149
- Future / 11

G

- Gadget / 56 57 134 135 162 191 192 193 198
- Gadget, Ticket Map / 134
- Gadgets / 56 58 191 198
- Gadgets, Dashboard / 191
- Gadgets, Referencing CorreLog Externally / 58
- Gateway / 200
- Gauge-alert-gadget / 193
- Gauge-loading-gadget / 193
- Gauge-ticket-gadget / 193 197
- GenDex Search Engine / 48
- Gendex / 47 48 76 78 162
- Gendexexe / 41 48
- Generate / 200
- Generation / 158 159
- Generators / 150
- Generators, Manually Running Report / 150
- Genrpt / 142
- Geo- / 176
- Gigabyte / 9 14
- Graph / 59 60 129 138 157
- Graph-message-rate / 58
- Graph-message-rate-gadget / 194
- Graph-message-rate-gadgetexe / 58
- Graph-thread-delta-gadget / 194

Graph-thread-rate / 194
Graph-thread-rate-gadget / 194
Graph Feature Summary / 59
Graph Scale Slider Bar / 60
Graphing / 157
Graphs / 59 78 174
Graphs Screen / 59
Group / 30 55 86 92 93 121 122 127 132 196
Group, Ticket Wizard / 127
Group name1 height / 56
Group name1 label / 55
Group name1 url / 55
Group name1 width / 55
Group name2 height / 56
Group name2 label / 56
Group name2 url / 56
Group name2 width / 56
Groups, Address / 92
Groups, Thread / 92
Gsip / 200 201
Guide / 13 23 81 96 97 98 158 176
Gzip / 47 75 79

H

Hardware / 141
Header / 179 183
Help / 96 198
Helpdesk ubat / 103
Helpdeskbat / 103
Helper / 53
Helper, DN Program / 53
Help... / 203
Hhmmss / 104 125
High / 8
Hint / 93
Hints / 30 93
Hipaa / 47 138 139 140
Hist / 149
History / 28 62 64 144
History, Catalog Message And Counters / 62
Home / 7 17 18 23 71 138 139 148 162 177 200 202
Hostname / 132
Hour / 130
Hourcount / 130 131
Hourly / 59 60

Hours / 130
Hourseverity / 130 131
Html / 57 58 138 152 161 172 173 195

I

Identification / 52
Identification, Managed Device / 52
Identifier / 52 176
Idle / 197
Iframe / 56 57 58 59
Ignore / 37
Impact / 126
Incident / 121
Incidents / 119
Include / 144 145 149
Incoming / 66
Incoming Message Filters / 66
Increasing / 94
Index / 176 200 205
Index, Alphabetical / 205
Info / 10 55 56 75 93 177 187 188
Inform / 114
Information / 9 25 46 47 52 53 54 55 126 129 163 188
Information, Catalog / 46
Information, Device Screen / 53
Information, From Viewpoint Of Theory / 188
Information, Log File Report / 47
Information, Reviewing Ticket / 25
Information, SNMP Device / 54
Insensitive / 87
Installation / 13 14 16 17 23 198
Installation, Basic CorreLog System Steps / 14
Installation, CorreLog / 13
Installer / 21
Installing / 17 18 105
Installing CorreLog Adapters / 18
Installing CorreLog WTS On Windows Platforms / 17
Instance / 95
Instances / 95 96 170
Integrity / 200
Interactive / 39 47 75
Interactive, CorreLog Usage / 47
Interested / 131
Interface / 105
Interface, Adding Correlation Action At Web / 105

Internal / 33 184
Interval / 130
Introduction / 5 7 8 5
Introduction To Correlation / 8
Introduction To Syslog Protocol / 7
Invalid / 143
Ipaddcnf / 46
Ipv4 / 20 52 53 90 147
Ipv6 / 20 52 200
Issues / 107
Issuing / 37
Item / 64 152 163 164 165 168
Item, Deleting Catalog / 64
Items / 142 151
Items, Excel Report Configuration And Settings / 142
Items, RSS Configuration And Settings / 151

J

Java / 17 100 101 126

K

Keep / 37 76 77 129 144 152 171
Kernel / 184
Keywd / 176
Keyword / 87 106 144 149 176
Keyworddata / 149
Keywords / 25 87

L

Label / 56
Large / 9
Last-message-gadget / 194
Last-message gadget / 194
Layout / 56 57 162 191
Layout, Dashboard Configuration And Files / 56
Ldap / 140 173
Le-notice / 52
Learn / 10 133
Line / 20 133
Links / 175
Linux / 44
List / 62 69 72 91 161 166 199
List, Windows Process / 199

List Macro Expressions / 91
Listenport / 44
Listenport-aux / 44
Listmode / 133
Lists / 91 92 169
Literal / 87
Live / 102
Load / 34
Load, Reducing Correlation And Message / 34
Local / 36 37 186
Local0 / 186
Local1 / 186
Local2 / 186
Local3 / 186
Local4 / 186
Local5 / 186
Local6 / 186
Local7 / 186
Local System Advisories / 36
Location / 55 70
Lock / 185
Lockout / 139
Log File Report Information / 47
Logfilebat / 101
Logins / 11 26 121 130 131 132 174 175
Loglocal / 45
Logons / 141
Logs / 7
Lookup / 176

M

Machines / 16 92
Macro / 91 98
Macro, List Expressions / 91
Macros / 91 92 98 106 169
Mail / 31 32 117 184
Main / 29 35
Maintexe / 41
Managed / 52
Managed Device Identification / 52
Management / 10 129 140 173
Management, Common Database Setup / 129
Manager / 6 8 16 20 40 41 45 78 114
Managers / 6
Manually / 123 150

Manually Adding Tickets / 123
Manually Running Report Generators / 150
Manuals / 103 158 159 177
Map-ticket-gadget / 134 195
Mask / 71 72
Master / 152 153
Match Expressions / 87
Matchaddress / 44
Matching / 51
Matching Severities / 51
Max-1 / 48
Max-50 / 48
Mcafee / 18 103 107 177
Meaning / 184 187
Meaning, Facility Codes And Their / 184
Message, Aux Files / 67
Message, Catalog And History Counters / 62
Message, Data Archiving Digest Functions / 76
Message, Excel Report Parsing Rules / 145
Message, Incoming Filters / 66
Message, Reducing Correlation And Load / 34
Message, Special Parsing Functions / 145
Message, Syslog Content / 182
Message Forwarding / 72
Message Overrides / 67
Messageaddress / 110
Messagedata / 102 109 110 149
Messagefacility / 110
Messages, Forwarding Syslog / 181
Messages, Related And Source Alert Definition / 123
Messages, Relaying / 111
Messages, Search Screen / 47
Messages, Sending E-Mail / 112
Messages, Sending SNMP Trap / 114
Messageseverity / 110
Messagetime / 110
Messenger / 102 117
Methods / 179
Mimicry / 180
Miscellaneous / 35
Mode / 130 131
Modes / 24
Modify / 28 36 148
Monitoring / 6 24
Monthly / 153
Ms-access / 154

Msg-01 / 73
Msg-04 / 73
Multi-tier / 74

N

Named / 148 156
Names / 53 70 71 72 156
Names, Auto-Masking User For Data Privacy / 71
Names, Auto-Overriding Device With Parsed Values / 70
Naming / 41
Naming, Syslog File Conventions / 41
Nesting / 88
Net-snmp / 114 167
Newest / 152 153
News / 137 184
Non-indexed / 25 63
None / 49 53 123 129 130
Normally / 77 108 116 123
Notation / 49 50
Notation, Advanced Search Term / 49
Notation, Search Term Examples / 50
Notes / 21 37 78 97 117 135 149 158 198
Notes, Additional / 21
Notes, Further On Excel Spreadsheets / 149
Notes, Section Summary And Additional / 37 78 158
Notice / 93 187
Notifications / 31 32
Notifications, Configuring SMTP / 31
Notifications, Configuring Ticket E-Mail And / 32
Numbers / 48
Numeric / 145 179

O

ODBC, Configuring DSN And Report Database / 154
ODBC, Support For Compliant Databases / 154
ODBC, Updating SQL Databases / 109
Odbc / 109 110 117 129 139 143 154 155 156 157 159 174 196
Office / 155
Online / 76
Opened / 29 122 123 124 125 129 170 171
Opening / 116
Opening, Automatically Tickets / 116
Openticketexe / 116 133 134 135
Operation / 24 74

- Operationally / 8
- Operations / 10
- Operator / 88
- Operators / 24 139
- Options / 31 195 198
- Oracle / 8 92 154
- Orchestrator / 103
- Organizing / 122
- Other Report Types / 157
- Other Syslog Configuration Files / 45
- Outlook / 31
- Output / 35
- Override / 10 69 70 72 166 167 188
- Overrides / 42 46 52 66 67 68 69 70 71 166 167 196
- Overrides, Message / 67
- Overriding / 53
- Overview / 41
- Overview, Syslog Process / 41

P

- Page / 48 205
- Parameter / 31 75
- Parameters / 116
- Parenthetical / 88
- Parse / 71 88 90 166 172 195
- Parse-thread-gadget / 195
- Parsed / 70 88
- Parsed, Auto-Overriding Device Names With Values / 70
- Parsing / 90 144 145
- Parsing, Excel Report Message Rules / 145
- Parsing, Special Message Functions / 145
- Partial / 87
- Path / 104
- Pattern / 82 95 192
- Patterns / 66 82 84 85 95 96 145 170
- Pci- / 138 139
- Pci-dss / 139 140
- Percentload / 130 131
- Perimeter / 140 172 176
- Perl / 17 100 101 106 126 186
- Permissions / 107
- Pick / 28
- Ping / 18 53 201
- Pivot / 157 158 159 174
- Pivot Reports / 158

- Platforms / 7 17
- Platforms, Installing CorreLog WTS On Windows / 17
- Platforms, Syslog Support On Win32 / 7
- Port / 43
- Port, Changing Syslog UDP Configuration / 43
- Ports / 14
- Positions / 145
- Post / 36 37 176
- Preference / 36
- Preferences / 30 37 56 191
- Prefix / 71 143 152 155 156 159
- Prefix, Unique Table / 155
- Prefs / 36 92 174 175
- Print / 157 158
- Printer / 184
- Priorities / 183
- Priority / 184
- Privacy / 71
- Privacy, Auto-Masking User Names For Data / 71
- Procedure / 142
- Process / 41 115 193 199
- Process, Syslog Overview / 41
- Process, Windows List / 199
- Processes / 40
- Profiles / 174
- Program / 7 8 20 53 76 101 103 106 107 108
- Program, CHKSC – Checking Archive Security Codes / 76
- Program, Creating New Action / 103
- Program, DN Helper / 53
- Program, Detailed Action Usage And Debugging / 107
- Program, Using Existing Action / 101
- Programs / 18 20
- Properties / 107
- Protection / 16 107
- Protocol / 7 179 180
- Protocol, BSD Syslog / 180
- Protocol, Introduction To Syslog / 7
- Protocol, Syslog / 179
- Proxy / 73
- Publish / 144 152
- Publishing / 153
- Publishing, RSS Security Considerations / 153

Q

- Queries / 138 172

Query / 25 35 47 76 129 137 138 147 171 172 195
Querying / 24
Queue / 201
Queued / 108 193
Quick / 13 23

R

RSS Configuration Items And Settings / 151
RSS Publishing Security Considerations / 153
RSS Report Distribution System / 151
Range / 105
Rate / 192
Readme / 146
Readmetxt / 146
Receiver / 6 19 111 115
Receivers / 73
Recent / 36 37 130
Recent, Auditing Advisories / 37
Reception / 8 43
Reception, Enabling Syslog TCP / 43
Records / 144
Recycle / 13 21
Redhat / 186
Reducing / 34
Reducing Correlation And Message Load / 34
Reduction / 9
Reference / 8 23 71 138 139 161 205
References / 96
References, Additional Correlation / 96
Referencing / 58
Referencing CorreLog Gadgets Externally / 58
Refining / 28 29
Regenerate / 27 29 65
Regenerating / 65
Related / 26 116 123 124 126
Related Messages And Source Alert Definition / 123
Relay / 72 115
Relaybat / 102 103 106 115
Relaying / 111
Relaying Messages / 111
Relocating / 43
Relocating Syslog Log Directory / 43
Remedy / 103 126
Repeat / 35
Report, Audit Types / 139

- Report, Configuring DSN And ODBC Database / 154
- Report, Excel Configuration Items And Settings / 142
- Report, Excel Message Parsing Rules / 145
- Report, Log File Information / 47
- Report, Manually Running Generators / 150
- Report, Other Types / 157
- Report, RSS Distribution System / 151
- Reporting / 6 9 10 77 137 158 165
- Reporting, CorreLog / 137
- Reports, Audit / 138
- Reports, Configuring E-Mail / 33
- Reports, Excel / 140
- Reports, Pivot / 158
- Reporttime / 148
- Reporttitle / 148
- Requirements / 14
- Resolution / 122 125
- Resources / 151
- Response / 53
- Results / 48
- Reviewing / 25
- Reviewing Ticket Information / 25
- Rfcs / 180
- Router / 17
- Routers / 17 121
- Rssexe / 153
- Rsyslog / 19 20 43 44
- Rule / 83 96
- Rule, Correlation Dependencies / 96
- Rule, Correlation s And Rule Systems / 83
- Rules / 83 85 90 144 145
- Rules, Configuration Of Correlation / 85
- Rules, Correlation And Rule Systems / 83
- Rules, Excel Report Message Parsing / 145
- Runmacroexe / 126
- Running / 150
- Running, Manually Report Generators / 150
- Runsql / 110 117
- Runsqlbat / 102 104 106 109 110 117
- Runsqlcnf / 109 117

S

- S-1-5-n-n / 176
- SMTP, Configuring Notifications / 31
- SNMP Device Information / 54

S address / 104
S facility / 104
S facility number / 104
S message / 104
S severity / 104
S severity number / 104
S summary / 104
S time / 104
S timestamp / 104
Same / 96
Save / 29 32 34 35 95
Saveas / 29
Saved / 138 172
Savenew / 33 95 162
Scale / 9 60
Scale, Graph Slider Bar / 60
Schedule / 19 43 153 175 201
Scheduler / 17 19 54
Score / 140 173
Screens / 56 66 75 100 161
Screens, Dashboard / 56
Scroll / 194 195
Search, Advanced Term Notation / 49
Search, Clearing Screen / 51
Search, GenDex Engine / 48
Search Messages Screen / 47
Search Term Notation Examples / 50
Search Terms / 48
Searching / 9 24 129
Search... / 163
Seconds / 67 130
Section Summary And Additional Notes / 37 78 158
Security / 5 6 24 76 153 175 176 184 189
Security, CHKSC Program – Checking Archive Codes / 76
Security, RSS Publishing Considerations / 153
Seed / 71
Select / 31 32
Send epobat / 103
Sending / 112 114
Sending E-Mail Messages / 112
Sending SNMP Trap Messages / 114
Sendlog / 8 75
Sendlogbat / 102 107 111 115 117
Sendmail / 32 100 175
Sendmailbat / 102 106 107 112 113 117
Sendpagebat / 112

Sendtrapbat / 102 107 114 117
Server, CorreLog Features / 8
Servers / 10 121
Service / 7 14 18 40 41 45
Services / 40
Session / 201
Sessions / 201
Setting / 77 127 130 144
Setting, Auto-Close Duplicate Tickets / 127
Settings / 142 151
Settings, Excel Report Configuration Items And / 142
Settings, RSS Configuration Items And / 151
Setup / 129
Setup, Common Management Database / 129
Severcnf / 46
Severities / 40 51 61 62 63 94 165 184 191 200
Severities-gadget / 195
Severities, Matching / 51
Severitydata / 149
Severitysummary / 149
Shield / 107
Siem / 73
Sigma / 13 21 133 154 157 174 175
Simple / 126 137 151 159
Site / 176 200
Size / 35
Slider / 60
Slider, Graph Scale Bar / 60
Sntp / 31 32 38 102 112 113 117 175 184
Snmpv2 / 114
Snmpv3 / 114
Source / 26 33 110 123 124 129 142 154 155 174 193
Source, Related Messages And Alert Definition / 123
Sources / 154
Space / 14
Span / 63 143 144 150
Special Message Parsing Functions / 145
Speed / 8
Spreadsheets / 149
Spreadsheets, Further Notes On Excel / 149
Sspi / 175
Stand / 6
Standard / 52 70
Start / 10 23 48 59
Start, CorreLog Fast / 10
Startpoint / 143 148 151 155 159

- Startup / 18 44 85
- State / 87
- Statistically / 93
- Statistics / 60 93 94
- Status / 122
- Step / 35
- Steps / 14 100
- Steps, Basic CorreLog System Installation / 14
- Steps, Detailed In Configuring User Action / 100
- Style / 179
- Subsequently / 71
- Successful / 141
- Suggest / 95 98
- Summary / 37 59 78 97 117 135 148 149 154 158
- Summary, Graph Feature / 59
- Summary, Section And Additional Notes / 37 78 158
- Sunday / 153
- Supply / 33
- Support / 7 70 74 154 158 177
- Support, Syslog On Win32 Platforms / 7
- Support For ODBC Compliant Databases / 154
- Svcexe / 78
- Sysldestport / 44
- Syslog- / 44
- Syslog-ng / 19 20 43 71
- Syslog, BSD Protocol / 180
- Syslog, Changing UDP Port Configuration / 43
- Syslog, Configuring On Network Computers / 16
- Syslog, Enabling TCP Reception / 43
- Syslog, Forwarding Messages / 181
- Syslog, Introduction To Protocol / 7
- Syslog, Other Configuration Files / 45
- Syslog, Relocating Log Directory / 43
- Syslog File Naming Conventions / 41
- Syslog Message Content / 182
- Syslog Process Overview / 41
- Syslog Protocol / 179
- Syslog Support On Win32 Platforms / 7
- Syslogdefender / 20
- Syslogexe / 15 40 43
- System, Basic CorreLog Installation Steps / 14
- System, Local Advisories / 36
- System, RSS Report Distribution / 151
- Systems / 83
- Systems, Correlation Rules And Rule / 83

T

- T assigned to / 125
- T comment / 125
- T date time / 125
- T exturl / 125
- T message / 125
- T related address / 126
- T related devdesc / 126
- T related devname / 126
- T related facility / 126
- T related message / 126
- T related severity / 126
- T related summary / 126
- T severity / 125
- T severity number / 125
- T state / 125
- T status / 124 125
- T summary / 125
- T time now / 125
- T tstamp / 125
- T uid / 124
- Table / 130 143 155 156 159 201
- Table, CMDB Definition / 130
- Table, Unique Prefix / 155
- Tabs / 53 55
- Task / 16 20 78 199 200 201 202
- Techniques / 73
- Techniques, Additional Forwarding / 73
- Template / 143 159
- Templatebat / 103 126
- Templates / 148 169
- Terabytes / 10
- Term / 49 50
- Term, Advanced Search Notation / 49
- Term, Search Notation Examples / 50
- Terms / 48
- Terms, Search / 48
- Testing / 106
- Testing Action / 106
- Text / 68 71 72 167 179
- Text-html-gadget / 195
- Their / 184 187
- Theory / 188
- Theory, From Viewpoint Of Information / 188
- Thread Groups / 92

Threading / 82
Threads, Defining New Correlation / 26
Threshold / 30 93 98 197
Ticket, Configuring E-Mail And Notifications / 32
Ticket, Reviewing Information / 25
Ticket Actions Facility / 124
Ticket Assignee / 121
Ticket Group Wizard / 127
Ticket Map Gadget / 134
Ticketbat / 116 134
Ticketing / 120 135
Tickets, Auto-Close Duplicate Setting / 127
Tickets, Automatically Opening / 116
Tickets, Closing And Deleting / 123
Tickets, CorreLog / 119
Tickets, Defining Correlation Alerts And / 30
Tickets, Editing / 122
Tickets, Manually Adding / 123
Tickets External URL / 132
Time / 42 48 53 55 87 105 109 149 197
Timeticks / 55
Title / 27 28 143
Tool / 7 8 17 18 20 44 75 176 202
Tools / 109 154 202
Top-alerts-gadget / 196
Top-devices-gadget / 196 197
Top-facilities-gadget / 196
Top-odbc-gadget / 196
Top-threads-gadget / 197
Top-ticket-users-gadget / 197
Top-users-gadget / 197
Tracking / 122
Transfer / 201
Translating / 71
Transport / 185
Trap / 44 111 114 202
Trap, Sending SNMP Messages / 114
Trapdestport / 44
Trigger / 82 83 84 85 87 95 97
Triggers / 8 9 82 83 84 85 91 95 96 98 168 169 170
Trimming / 9
Tunnel / 73 103 115
Tunnelbat / 103 115
Tunnelexe / 44
Tunneling / 202
Types / 139 157

Types, Audit Report / 139
Types, Other Report / 157
Typically / 17

U

Unattended / 6 24
Uninstall / 21
Uninstalling / 20
Uninstalling Entire CorreLog Framework / 20
Unique / 143 155 156
Unique Table Prefix / 155
Update / 121 123 124 130
Updating / 109
Updating SQL ODBC Databases / 109
Urls / 53 55 56 174
Usage / 23 47 107
Usage, CorreLog Basic / 23
Usage, CorreLog Interactive / 47
Usage, Detailed Action Program And Debugging / 107
User, Auto-Masking Names For Data Privacy / 71
User, Creating Defined Facility Codes / 69
User, Deleting Defined Facility Codes / 69
User, Detailed Steps In Configuring Action / 100
Username / 129 130 131
Usersummary / 148
Using Existing Action Program / 101
Utilities / 11
Uucp / 181

V

Value / 71
Values / 70
Values, Auto-Overriding Device Names With Parsed / 70
Variables / 103
Variables, Correlation Actions Environmental / 103
Versions / 152
Viewer / 63 64 138
Viewers / 139
Viewing / 10
Viewpoint / 188
Viewpoint, From Of Information Theory / 188
Virus / 14 16 107
Visit / 203
Vista / 8 17 20

Visual / 150

W

Warning / 93 187

Watch / 197

Watch-device-gadget / 197

Wbem / 179

Weekly / 153

Wildcard / 87

Win200x / 8

Win32 / 7 11 17

Win32, Syslog Support On Platforms / 7

Windows, Installing CorreLog WTS On Platforms / 17

Windows Process List / 199

Winmsgrbat / 102

Winzip / 13 15 18

Wizard / 27 30 32 33 85 86 87 100 127 135 162

Wizard, Ticket Group / 127

Wizards / 85

Words / 145

Working / 107



CorreLog, Inc.

<http://www.CorreLog.com>

<mailto:support@CorreLog.com>

Copyright © 2008 - 2014. All rights reserved.